

The Costs of Environmental Regulation in a Concentrated Industry*

Stephen P. Ryan[†]

July 1, 2011

Abstract

The typical cost analysis of an environmental regulation consists of an engineering estimate of the compliance costs. In industries where fixed costs are an important determinant of market structure this static analysis ignores the dynamic effects of the regulation on entry, investment, and market power. I evaluate the welfare costs of the 1990 Amendments to the Clean Air Act on the US Portland cement industry, accounting for these effects through a dynamic model of oligopoly in the tradition of Ericson and Pakes (1995). Using the two-step estimator of Bajari, Benkard, and Levin (2007), I recover the entire cost structure of the industry, including the distributions of sunk entry costs and capacity adjustment costs. My primary finding is that the Amendments have significantly increased the sunk cost of entry, leading to a loss of between \$810M and \$3.2B in product market surplus. A static analysis misses the welfare penalty on consumers, and obtains the wrong sign of the welfare effects on incumbent firms.

*This is a revised version of my job market paper. I would like to especially thank Pat Bajari for guidance and support. I have also benefited from conversations with Tom Ahn, Arie Beresteanu, Jane Cooley, Paul Ellickson, Han Hong, Shanjun Li, Chris Timmins, Justin Trogon, and numerous seminar participants. Comments from three referees and the editor have vastly improved the paper. All remaining errors are my own.

[†]MIT Department of Economics and NBER.

1 Introduction

In the United States, the Environmental Protection Agency (EPA) is responsible for setting and enforcing regulations broadly consistent with national environmental policies, such as the Clean Air Act (CAA). The CAA gives the EPA a mandate to regulate the emissions of airborne pollutants such as ozone, sulfur dioxide (SO_2), and nitrogen oxides (NO_x), in the hopes of producing a healthier atmosphere. The Clean Air Act and its subsequent Amendments require the Agency to assess the costs and benefits of a regulation before promulgating policy. The cost analysis is typically an engineering estimate of the expenditures on control and monitoring equipment necessary to bring a plant into compliance with the new regulations. However, this type of cost analysis misses most of the relevant economic costs in concentrated industries, in which sunk costs of entry and costly investment are important determinants of market structure. Shifts in the costs of entry and investment can lead to markets with fewer firms and lower production. The resulting increase in market concentration can have far-reaching welfare costs beyond the initial costs of compliance. This is a particularly acute problem for environmental regulators, as many of the largest polluting industries are also highly concentrated.

In this paper, I measure the welfare costs of the 1990 Clean Air Act Amendments on the US Portland cement industry, explicitly accounting for the dynamic effects resulting from a change in the cost structure. Portland cement is the binding material in concrete, a primary construction material found in numerous applications, such as buildings and highways. The industry is typical of many heavy industries, consuming large quantities of raw materials and generating significant amounts of pollution byproducts. It is a frequent target of environmental activists and has been heavily regulated under the Clean Air Act. In 1990, Congress passed Amendments to the Clean Air Act, adding new categories of regulated emissions and requiring plants to undergo an environmental certification process. It has been the most comprehensive and important new environmental regulation affecting this industry in the last three decades since the original Clean Air Act.

My strategy for evaluating the effects of the Amendments on this industry proceeds in three distinct steps. First, I pose a theoretical model of the cement industry, where oligopolists make optimal decisions over entry, exit, production, and investment given the strategies of their competitors. Second, using a unique panel data set covering two

decades of the Portland cement industry, I recover parameters which are consistent with the underlying model. Third, I use the theoretical model to simulate economic environments with the cost structures recovered before and after the Amendments. By comparing the predictions of the model under these different cost structures, I can calculate the changes to a number of quantities relevant to policymakers, such as producer profits and consumer surplus, that are the result of the regulation.

The backbone of my analysis is a fully dynamic model of oligopoly in the tradition of Maskin and Tirole (1988) and Ericson and Pakes (1995). I model the interaction of firms in spatially-segregated regional markets where firms are differentiated by production capacity. Firms are capacity constrained and compete over quantities in homogeneous good markets. Markets evolve as firms enter, exit, and adjust their capacities in response to variation in the economic environment. I incorporate sunk costs of entry, fixed and variable costs of capacity adjustment, and a fixed cost of exiting the industry. I assume that firms optimize their behavior conditional only on the current state vector and their private shocks, which results in a Markov-perfect Nash equilibrium (MPNE).

The MPNE of the model leads to structural requirements on firm behavior which can be used as the basis of an estimator of the underlying primitives. As Benkard (2004) illustrates, the impediment to using these types of models for empirical work has been the computational burden of solving for the MPNE, which makes nested fixed-point estimators in the tradition of Rust (1987) impractical. However, a series of recent papers has built on the insights of Hotz, Miller, Sanders, and Smith (1994) to circumvent this problem using a two-step approach, in which it is possible to estimate the dynamic model without solving for the equilibrium even once.¹

In these two-step estimators, the econometrician first simply describes *what* the firms do at every state, and then imposes equilibrium restrictions from an underlying model to explain *why* the firms behave as they do. This approach circumvents the need to compute the equilibrium to the model as part of the estimation process. Following the simulation-based minimum distance estimator proposed by Bajari, Benkard, and Levin (2007), I estimate both the distribution of fixed adjustment costs and variable adjustment costs, the distribution of scrap values associated with exiting the market,

¹Representative papers in this literature include Bajari, Benkard, and Levin (2007), Aguirregabiria and Mira (2007), Pakes, Ostrovsky, and Berry (2007), and Pesendorfer and Schmidt-Dengler (2008).

and the distribution of entry costs. I recover these parameters before and after the 1990 Amendments in order to evaluate the changes in the underlying cost structure induced by the Amendments. My primary empirical finding is that the Amendments led to a marked increase in the expected entry costs that a firm has to pay to enter the industry, while the other cost parameters remained unchanged.

After recovering estimates for the underlying model primitives, I solve for the MPNE of the theoretical model. I then simulate the model to calculate expected producer and consumer welfare, the number and size of firms, and the distribution of costs across incumbents and potential entrants before and after the regulations. I do not consider any benefits accruing to consumers due to reduced emissions, given the difficulty of quantifying the amount of emissions from cement plants and their associated damages.² I find that overall product market welfare has decreased between \$820M and \$3.2B as a result of the Amendments, due to an increase in the average sunk cost of entry. More importantly, as my estimates show the costs of production have not changed significantly after the regulations, the welfare effect on producers depends critically on whether or not the firm is an incumbent. While potential entrants suffer welfare losses as the result of paying higher entry costs, incumbent firms benefit from increased market power due to reduced competition. A static analysis of this industry would preclude changes in barriers to entry, and would obtain the wrong sign for the welfare costs of the Amendments on incumbent firms while understating overall welfare costs by at least \$300M.

2 Industry Background

Portland cement is a fine mineral dust with useful binding properties that make it the key ingredient of concrete, a pourable material composed of water, cement, and aggregate such as sand and stone. The concrete is then used as a fill material, such as in highways and buildings, and in finished products like concrete blocks.

The production of cement requires two commodities in enormous quantities: limestone and heat. The limestone is usually obtained from a quarry located at the production site. Large chunks of limestone are pulverized before being sent to the centerpiece of cement operations: an enormous rotating kiln furnace. These kilns are the largest moving pieces of industrial equipment in the world; they range in length

²I am continuing this line of research in Fowlie, Reguant, and Ryan (2010).

from 450 to 1000 feet and have diameters of over 15 feet. The chemical process of converting limestone into cement requires temperatures equal to a third of those found on the surface of the sun, so one end of the kiln is heated with an intense flame produced by burning fossil fuels. These high energy requirements are what lead the cement industry, a tiny part of the US economy at under \$10B a year in revenues, to have a large role in the environmental debate over emissions. Furthermore, the chemistry of the production of cement liberates carbon dioxide as a byproduct, which means that the production of cement is a major contributor of greenhouse gases globally.

Cement is a difficult commodity to store, as it will gradually absorb water out of the air, rendering it useless. As a result, producers and distributors do not maintain large stocks. Also, I treat cement as a homogeneous good since producers in the United States adhere to the American Society for Testing and Materials *Specification for Portland Cement*. Cement's use as a construction material means that producers are held to strict conformity with these specifications.

Cement is difficult to store for long periods of time, and therefore transportation costs are the most significant factors in determining Portland cement markets. Average transportation costs reported by US producers for shipments within 50 miles of the plant were \$5.79 per ton. These costs increased to \$9.86 per ton for shipments within 51-100 miles, \$14.53 per ton for 101-200 miles, and to \$18.86 per ton for 201-300 miles. For shipments that are 500 miles or more from the plant, transportation costs increased to \$25.85 per ton.³ These high costs, in conjunction with cement's low unit value, are the principal reasons the majority of cement is shipped locally.⁴

In 2000, the domestic Portland cement industry consisted of 116 plants in 37 states, operated by one government agency and approximately 40 firms. The industry produced 86 million tons of Portland cement with a raw value of approximately \$8.7 billion; most of this was used to make concrete, with a final value greater than \$35 billion. Domestic cement production accounted for the vast majority of the cement used in the United States. According to the US Geological Survey (2001), about 73 percent of cement sales were to ready-mixed concrete manufacturers, 12 percent to concrete product producers, 8 percent to contractors, 5 percent to building materials dealers, and 2 percent for other uses. Cement expenditures in construction projects

³These figures are taken from American University's Trade and Environment Database (TED) case study on Cemex.

⁴Jans and Rosenbaum (1997) quote a Census of Transportation report stating that 82.5 percent of cement was shipped under 200 miles, with 99.8 percent being shipped under 500 miles.

Table 1: Cement Industry Summary Statistics

Year	Production	Imports	Consumption	Price	Capacity	Capacity Per Kiln
1980	68,242	3,035	70,173	111.90	89,561	239
1981	65,054	2,514	66,092	103.70	93,203	267
1982	57,475	2,231	59,572	95.76	89,770	287
1983	63,884	2,960	65,838	91.01	92,052	292
1984	70,488	6,016	76,186	89.70	91,048	297
1985	70,665	8,939	78,836	84.71	88,600	305
1986	71,473	11,201	82,837	81.48	87,341	305
1987	70,940	12,753	84,204	78.07	86,709	314
1988	69,733	14,124	83,851	75.50	86,959	327
1989	70,025	12,697	82,414	72.04	84,515	337
1990	69,954	10,344	80,964	69.02	83,955	345
1991	66,755	6,548	71,800	66.37	84,471	352
1992	69,585	4,582	76,169	64.25	85,079	357
1993	73,807	5,532	79,701	63.58	84,869	363
1994	77,948	9,074	86,476	68.06	85,345	364
1995	76,906	10,969	86,003	72.56	86,285	367
1996	79,266	11,565	90,355	73.64	85,687	376
1997	82,582	14,523	96,018	74.60	86,465	383
1998	83,931	19,878	103,457	76.45	87,763	393

Summary statistics for the Portland cement industry 1980-1998. The data is from [Historical Statistics for Mineral and Materials Commodities in the United States](#), an online publication of the US Geological Survey. The units on quantities are thousands of metric tons, while prices are denoted in 1998 constant dollars.

are usually on the order of less than 2 percent of total outlays.

Table 1 reports summary statistics for the industry over the period 1980-1998. One point of interest is that capacity utilization rates have risen since the passage of the Amendments. Production has increased while overall productive capacity has remained relatively steady. Imports grew as the production of domestic cement reached its maximum level, and firms chose to import instead of build new production facilities.⁵ The industry has become slightly more concentrated over time. According to the Economic Census (2002), which collects extensive information on American industry every five years, the national four-firm concentration ratio was 38.7 in 2002, 33.5 in 1997, 28 in 1992, 28 in 1987, 31 in 1982, and 24 in 1977. However, these numbers mask the regional variation in cement concentration, as national market share may not be representative of competitive conditions in any given geographic region, due to the local nature of the cement industry.

The effects of imports on domestic producers are difficult to quantify due to the idiosyncracies associated with distributing cement from waterborne sources. For most markets, the economic impact is small and indirect, as few regions have the infrastructure and geography to profitably exploit the availability of imports. An examination of the import data provided in the USGS reports indicates that cement imports vary

⁵Cement imports come primarily from Canada, China, Korea, Thailand, Spain, and Venezuela. Asian sources have become the dominant source of cement imports, with Thailand becoming the single-largest exporter in 2000.

widely across markets and across time. Imported cement is actually shipped as clinker, the unground precursor of cement. In order to turn this raw material into cement, the importer must have a grinder and a supply of gypsum. Additionally, domestic cement producers have been highly successful in preventing large-scale imports through trade tariffs. For example, producers in states bordering the Gulf of Mexico have been successful in getting anti-dumping tariffs passed against imports from Mexico. This has limited the ability of importers to achieve greater penetration of local cement markets in these states. In large part, the response of potential importers has been to circumvent the tariffs through the acquisition of domestic facilities. In markets where imports do play a significant long-run role in the domestic market, such as around the Great Lakes region, I model this as a permanent shifter in the demand curve for domestically-produced cement.⁶

There have been two major regulatory events of interest to the Portland cement industry in the last 30 years: the Clean Air Act of 1970 and its subsequent Amendments in 1990.⁷ The stated purpose of the Clean Air Act was to “protect and enhance the quality of the Nation’s air resources so as to promote the public health and welfare and productive capacity of its population.” To this end, Congress empowered the EPA to set and enforce environmental regulations governing the emission of airborne pollutants.

In 1990, Congress passed the Amendments to the Clean Air Act, which defined new categories of regulated pollutants and required major polluters to obtain a permit for operation. These Amendments mandated new monitoring, reporting, and emission requirements for the cement industry. The Amendments created a new class of emission restrictions governing hazardous air pollutants and volatile organic compounds. One key identifying feature of this legislation is that EPA did not promulgate

⁶I am exploring the interactions of environmental regulation and cement imports in ongoing work. The potential for emissions leakage, e.g. Fowlie (2009), to undo domestic regulations is an important question in the environmental literature, and could mitigate any benefits from the Amendments.

⁷There have been other substantial changes to environment policy during the time period as well. One such change was the New Source Review (NSR) requirements instituted in the 1977 Amendments to the Clean Air Act. The NSR requires firms to obtain costly permits before substantially modifying older capital equipment or building new plants, and is believed by the EPA (2001) to have changed the pattern of investment in other industries, such as power plants and refineries. The NSR also likely created barriers to entry in this industry, as new entrants may have faced higher capital costs than incumbents, and plants located in lower pollution areas may face different emissions requirements than those in higher pollution areas. Since the NSR requirements did not change over the sample period that I am examining, I do not explicitly model the differential effects of the NSR on firms in the cement industry.

final requirements for these new pollutants for 12 years. Therefore, there were no direct changes to firms' variable costs as a result of the Amendments, as they did not require the firms to adhere to any new emissions standards.⁸

There were two components of the legislation that began to bind immediately. Under Title V of the Amendments, all firms emitting significant quantities of pollutants had to apply for operating permits. The permits require regular reporting on emissions, which necessitate the installation and maintenance of new monitoring equipment. The Amendments also required firms to draw up formal plans for compliance and undergo certification testing. Industry estimates for the costs of compliance with these operating permits is on the order of five to ten million dollars. By 1996, virtually all cement plants had applied for their permits, which they are required to renew every five years. The EPA estimated that these certification costs would not exceed \$5M per establishment.

The second aspect of the Amendments which is critical to understanding their welfare implications is that they required brand new plants to undergo an additional, rigorous environmental certification and testing procedure. These additional fixed costs involved potential entrants contracting with environmental engineering firms to produce reports on their impact on local air and water quality as a result of the construction and operation of a new plant. Industry sources estimate that these costs would add approximately \$5M to \$10M to the cost of building a greenfield facility. It is this change to the sunk costs of entry which is going to drive many of the results that I find below.

3 Data Sources

I collect data on the Portland cement industry from 1980 to 1999 using a number of different sources. I require market-level data on prices and quantities to estimate the demand curve for cement. The US Geological Survey (USGS) collects establishment-level data for all the Portland cement producers in the US and publishes the results

⁸To the best of my knowledge, as of 2009 no firm has made any changes to its production process as a result of the Amendments, due in part to legal opposition from the Portland Cement Association. Firms may also reasonably anticipate that changes to their marginal costs may ultimately be close to zero, as either they will be grandfathered into the legislation or the EPA may give pollution credits in return for adopting lower emissions standards.

Table 2: Summary Statistics

Variable	Minimum	Mean	Maximum	Standard Deviation
Market-level Supply and Demand Data				
Quantity	186	2,835.84	10,262	1,565.34
Price	36.68	67.46	138.99	13.68
Plants In Market	1	4.75	20	1.94
Skilled Wage	20.14	31.72	44.34	4.33
Coal Price	15.88	26.64	42.33	8.13
Electricity Price	4.23	5.68	7.6	1.01
Natural Gas Price	3.7	6.21	24.3	2.21
Population	689,584	10,224,352	33,145,121	7,416,485
Plant-level Production Data				
Quantity	177	699	2348	335
Capacity	196	797	2678	386
Plant-level Investment				
Capacity Investment	-728	2.19	1,140	77.60

Demand data are from annual volumes of the USGS's Mineral Yearbook, 1980-1981 to 1998-1999. There are 517 observations in 27 regional markets. Quantities and capacities are denominated in thousands of tons per year, while price is denoted in dollars per ton. Labor wages are denoted in dollars per hour for skilled manufacturing workers, and taken from County Business Patterns. Population is the total populations of the states covered by a regional market. The units are dollars per ton for coal, dollars per kilowatt hour for electricity, and dollars per thousand cubic feet for gas. All prices are adjusted to 1996 constant dollars. The data on production and capacity are taken from the Portland Cement Association's annual Plant Information Summary, with full coverage from 1980 to 1998.

in their annual Minerals Yearbook.⁹ The USGS aggregates establishment-level data into regional markets to protect the confidentiality of the respondents. The Minerals Yearbook contains the number of plants in each market and the quantity and prices of shipped cement.¹⁰

This data is the source of market definitions that I use through the remainder of the analysis. The USGS examines the set of firms which compete with one another, and aggregates their locations together into a market. While this is clearly an imperfect measure of market definition, since prices in neighboring markets almost surely have influence on the prices within a given market, it is roughly consistent with the idea that cement is very expensive to ship long distances. If one were to draw 100 mile circles around each cement plant in the United States, the resulting areas of significant overlap look much like the market definitions from the USGS data.

I collect data on electricity prices, coal prices, natural gas prices, and manufacturing wages to use as instruments in the demand curve estimation. The data for fuel

⁹The Bureau of Mines had this responsibility prior to merging with the USGS in the 1990s. The data was collected by a mail survey, with a telephone follow-up to non-respondents. Typically the total coverage of the industry exceeded 90 percent; in some years, 100 percent response was indicated. The USGS attempted to fill in missing observations with data from other sources.

¹⁰There is occasional irregular censoring of data to ensure the confidentiality of individual companies, although this affects only a small number of observations representing a low percentage of overall quantity. Usually the USGS merges a censored region into a larger region in subsequent years to facilitate complete reporting.

and electricity prices are from the US Department of Energy’s Energy Information Administration.¹¹ Natural gas and electricity prices are reported at the state level from 1981 to 1999. Coal prices are only available in a full series over that time span at the national average level. I impute skilled manufacturing wages at the state level from the US Census Bureau’s County Business Patterns. All prices are adjusted to 1996 constant dollars. I also collected market-level data on population and housing permits from the US Census Bureau.¹²

Table 2 shows summary statistics for the demand data. Most markets are characterized by a small number of firms, with the median market contested by four firms. The size of the markets varies greatly across the sample: the smallest market is two percent of the size of the largest market. Price also varies substantially across markets, with Alaska and Hawaii generally being the most expensive markets.

Data on the plant-level capacities and production quantities are from the Portland Cement Association’s annual Plant Information Summary (PIS) and cover 1980 to 1998. For each establishment in the US, the PIS reports daily and annual plant capacities. I interpret the daily capacity to be a boilerplate rating, determined by the manufacturer of the kiln at the time of its manufacture, of how much the kiln produces in a given 24-hour period of operation. A critical assumption that I make is that I interpret the number listed as yearly capacity as representing how much cement that plant actually produced in that year.¹³

I emphasize, however, that production quantity is not exogenously set as a fixed percentage of the theoretical maximum capacity, as firms still choose how long to operate their kilns before performing maintenance. Given that firms are at the edge of their maximum productive capacity during the sample period, capacity choice is

¹¹<http://www.eia.doe.gov>.

¹²See <http://www.census.gov/popest/states/> and <http://www.census.gov/const/www/permitsindex.html>, respectively. Housing permits are broken out by size of the dwelling, in residential units, up to structures with five or more units.

¹³This assumption is supported by the fact that plants operate continuously in runs lasting most of the year except for a maintenance period, generally a month in duration, in which the plant produces nothing. If the firms are assumed to run at perfect efficiency on the days which they operate, then the boilerplate rating multiplied by the length of a year gives the theoretical maximum that a plant could have produced. These boilerplate ratings typically do not change from year to year. On the other hand, the yearly capacity numbers never achieve this bound and fluctuate from year to year. Additionally, the yearly numbers approximate the market-level quantities reported in the USGS data, which was collected through a confidential survey of cement manufacturers. Therefore, I interpret the reported annual capacity of the kiln to be the amount of cement that it actually produced in that year.

clearly the most important strategic decision firms have to make, but it should be emphasized that they still face a tradeoff between production and maintenance. The last two rows of Table 2 give the summary statistics for production and capacity levels.

A key empirical fact of this industry is that most firms do not make adjustments to their capacity in most periods. The modal adjustment is zero, with a mean of just 2.9 thousand tons per year (TPY). While there is some noise in the data, it is clear that most firms have relatively steady levels of capacity over time, with infrequent discrete adjustments. In addition to capacity investment, there are jumps in market-level capacity due to entry and exit.

To match the market-level demand data to the establishment data from the PIS, I combine some of the markets in the USGS data to form continuously-reported metamarkets. I then group all the plants into the appropriate metamarkets for every year of establishment data. The production data consists of an unbalanced panel of 2,233 observations.

4 Model

My theoretical model of the cement industry builds on Maskin and Tirole (1988) and Ericson and Pakes (1995), who provide an elegant theoretical framework of industry dynamics.¹⁴

The basic building block of the model is a regional cement market. Each market is fully described by the $\bar{N} \times 1$ state vector, s_t , where s_{it} is the capacity of the i -th firm at time t , and \bar{N} is an exogenously-imposed maximal number of active firms. Firms with zero capacity are considered to be potential entrants. Time is discrete and unbounded, and firms discount the future at rate $\beta = 0.9$. Each decision period is one year. In each period, the sequence of events unfolds as follows: first, incumbent firms receive a private draw from the distribution of scrap values, and decide whether or not to exit the industry. Potential entrants receive a private draw from the distribution of both investment and entry costs, while incumbents who have decided not to exit receive private draws on the fixed costs of investment and divestment. All firms

¹⁴My model is similar to several other applications of the Ericson-Pakes framework; see, for example, Fershtman and Pakes (2000), Gowrisankaran and Town (1997), Besanko and Doraszelski (2004), Doraszelski and Satterthwaite (2010), and Benkard (2004).

then simultaneously make entry and investment decisions. Third, incumbent firms compete over quantities in the product market. Finally, firms enter and exit, and investments mature. I assume that firms who decide to exit produce in this period before leaving the market, and that adjustments in capacity take one period to realize. I also assume that each firm operates independently across markets.¹⁵

Firms obtain revenues from the product market and incur costs from production, entry, exit, and investment. Firms compete in quantities in a homogeneous goods product market. In each market m , firms face a constant elasticity of demand curve:

$$\ln Q_m(\alpha) = \alpha_{0m} + \alpha_1 \ln P_m, \quad (1)$$

where Q_m is the aggregate market quantity, P_m is price, α_{0m} is a market-specific intercept, and α_1 is the elasticity of demand. The cost of output, q_i , is given by the following function:

$$C_i(q_i; \delta) = \delta_0 + \delta_1 q_i + \delta_2 1(q_i > \nu s_i)(q_i - \nu s_i)^2. \quad (2)$$

Fixed costs of production are given by δ_0 . Variable production costs consist of two parts: a constant marginal cost, δ_1 , and an increasing function that binds as quantity approaches the capacity constraint. I assume that costs increase as the square of the percentage of capacity utilization, and parameterize both the penalty, δ_2 , and the threshold at which the costs bind, ν . This second term, which gives the cost function a “hockey stick” shape common in the electricity generation industry, accounts for the increasing costs associated with operating near maximum capacity, as firms have to cut into maintenance time in order to expand production beyond utilization level ν . I assume that firms play a capacity-constrained Cournot quantity game in each period.¹⁶ I denote the profits accruing from the product market by $\bar{\pi}_i(s; \alpha, \delta)$.

Firms can change their capacity through costly adjustments, x_i . The cost function

¹⁵This assumption explicitly rules out more general behavior, such as multimarket contact as considered in Bernheim and Whinston (1990) and Jans and Rosenbaum (1997).

¹⁶In the presence of fixed operation costs the product market may have multiple equilibria, as some firms may prefer to not operate given the outputs of their competitors. However, if all firms produce positive quantities then the equilibrium vector of production is unique, as the best-response curves are downward-sloping.

associated with these activities is given by:

$$\Gamma(x_i; \gamma) = 1(x_i > 0)(\gamma_{i1} + \gamma_2 x_i + \gamma_3 x_i^2) + 1(x_i < 0)(\gamma_{i4} + \gamma_5 x_i + \gamma_6 x_i^2). \quad (3)$$

Firms face both fixed and variable adjustment costs that vary separately for positive and negative changes. Fixed costs capture the idea that firms may have to face significant setup costs, such as obtaining permits or constructing support facilities, that accrue regardless of the size of the kiln. Fixed positive investment costs are drawn each period from the common distribution F_γ , which is distributed normally with mean μ_γ^+ and standard deviation σ_γ^+ , and are private information to the firm. Divestment sunk costs may be positive as the firm may encounter costs in order to shut down the kiln and dispose of related materials and components. On the other hand, firms may receive revenues from selling off their infrastructure, either directly to other firms or as scrap metal.¹⁷ These costs are also private information, and are drawn each period from the common distribution G_γ , which is distributed normally with mean μ_γ^- and standard deviation σ_γ^- .

Firms face fixed costs unrelated to production, given by $\Phi_i(a)$, which vary depending on their current status and chosen action, a_i :

$$\Phi_i(a_i; \kappa_i, \phi_i) = \begin{cases} -\kappa_i & \text{if the firm is a new entrant,} \\ \phi_i & \text{if the firm exits the market.} \end{cases} \quad (4)$$

Firms that enter the market pay a fixed cost of entry, κ_i , which is private information and drawn from the common distribution of entry costs, F_κ . Firms exiting the market receive a payment of ϕ_i , which represents net proceeds from shuttering a plant, such as selling off the land and paying for an environmental cleanup. This value may be positive or negative, depending on the magnitude of these opposing payments. The scrap value is private information, drawn anew each period from the common distribution, F_ϕ . Denote the activation status of the firm in the next period as χ_i , where $\chi_i = 1$ if the firm will be active next period, whether as a new entrant or a continuing incumbent, and $\chi_i = 0$ otherwise. All of the shocks that firms receive each

¹⁷One online example of a used market for cement equipment is www.usedcementequipment.com. While the prices of used equipment may be low, or even nominally zero, transportation and cleanup costs are typically high, occasionally into the millions of dollars depending on the size and type of equipment.

period are mutually independent.

Collecting the costs and revenues from a firm’s various activities, the per-period payoff function is:

$$\pi_i(s, a; \alpha, \delta, \gamma_i, \kappa_i, \phi_i) = \bar{\pi}_i(s; \alpha, \delta) - \Gamma(x_i; \gamma_i) + \Phi_i(a_i; \kappa_i, \phi_i). \quad (5)$$

For the sake of brevity, I henceforth denote the vector of parameters in Equation 5 by θ .

4.1 Transitions Between States

To close the model it is necessary to specify how transitions occur between states as firms engage in investment, entry, and exit. I assume that changes to the state vector through entry, exit, and investment take one period to occur and are deterministic. The first part is a standard assumption in discrete time models, and is intended to capture the idea that it takes time to make changes to physical infrastructure of a cement plant. The second part abstracts away from depreciation, which does not appear to be a significant concern in the cement industry, and uncertainty in the time to build new capacity.¹⁸

4.2 Equilibrium

In each time period, firm i makes entry, exit, production, and investment decisions, collectively denoted by a_i . Since the full set of dynamic Nash equilibria is unbounded and complex, I restrict the firms’ strategies to be anonymous, symmetric, and Markovian, meaning firms only condition on the current state vector and their private shocks when making decisions, as in Maskin and Tirole (1988) and Ericson and Pakes (1995).

Each firm’s strategy, $\sigma_i(s, \epsilon_i)$, is a mapping from states and shocks to actions:

$$\sigma_i : (s, \epsilon_i) \rightarrow a_i, \quad (6)$$

where ϵ_i represents the firm’s private information about the cost of entry, exit, investment, and divestment. In the context of the present model, $\sigma_i(s)$ is a set of policy

¹⁸It is conceptually straightforward to add uncertainty over time-to-build in the model, but assuming deterministic transitions greatly reduces the computational complexity of solving for the model’s equilibrium.

functions which describes a firm's production, investment, entry, and exit behavior as a function of the present state vector. In a Markovian setting, with an infinite horizon, bounded payoffs, and a discount factor less than unity, the value function for an incumbent at the time of the exit decision is:

$$\begin{aligned}
V_i(s; \sigma(s), \theta, \epsilon_i) = & \bar{\pi}_i(s; \theta) + \max \left\{ \phi_i, E_{\epsilon_i} \left\{ \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) dP(s'; s, \sigma(s)) \right. \right. \\
& + \max_{x_i^* > 0} \left[-\gamma_{i1} - \gamma_2 x_i^* - \gamma_3 x_i^{*2} + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right], \\
& \left. \left. \max_{x_i^* < 0} \left[-\gamma_{i4} - \gamma_5 x_i^* - \gamma_6 x_i^{*2} + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right] \right\} \right\}, \tag{7}
\end{aligned}$$

where θ is the vector of payoff-relevant parameters, E_{ϵ_i} is the expectation with respect to the distributions of shocks, and $P(s'; \sigma(s), s)$ is the conditional probability distribution over future state s' , given the current state, s , and the vector of strategies, $\sigma(s)$.

Potential entrants must weigh the benefits of entering at an optimally-chosen level of capacity against their draws of investment and entry costs. Firms only enter when the sum of these draws is sufficiently low. I assume that potential entrants are short-lived; if they do not enter in this period they disappear and take a payoff of zero forever, never entering in the future.¹⁹ Potential entrants are also restricted to make positive investments; firms cannot “enter” the market at zero capacity and wait for a sufficiently low draw of investment costs before building a plant. The value function for potential entrants is:

$$\begin{aligned}
V_i^e(s; \sigma(s), \theta, \epsilon_i) = & \max \{ 0, \\
& \max_{x_i^* > 0} \left[-\gamma_{1i} - \gamma_2 x_i^* - \gamma_3 x_i^{*2} + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right] - \kappa_i \}. \tag{8}
\end{aligned}$$

Markov perfect Nash equilibrium (MPNE) requires each firm's strategy profile to

¹⁹This assumption is for computational convenience, as otherwise one would have to solve an optimal waiting problem for the potential entrants. See Ryan and Tucker (2010) for an example of such an optimal waiting problem.

be optimal given the strategy profiles of its competitors:

$$V_i(s; \sigma_i^*(s), \sigma_{-i}(s), \theta, \epsilon_i) \geq V_i(s; \tilde{\sigma}_i(s), \sigma_{-i}(s), \theta, \epsilon_i), \quad (9)$$

for all s , ϵ_i , and all possible alternative strategies, $\tilde{\sigma}_i(s)$. As I work with the expected value functions below, I note that the MPNE requirement also holds after integrating out firms' private information: $E_{\epsilon_i} V_i(s; \sigma_i^*(s), \sigma_{-i}(s), \theta, \epsilon_i) \geq E_{\epsilon_i} V_i(s; \tilde{\sigma}_i(s), \sigma_{-i}(s), \theta, \epsilon_i)$. Doraszelski and Satterthwaite (2010) discuss the existence of pure strategy equilibria in settings similar to the one considered here. The introduction of private information over the discrete actions guarantees that at least one pure strategy equilibrium exists, as the best-response curves are continuous. However, there are no guarantees that the equilibrium is unique, a concern I discuss next in the context of my empirical approach.

5 Empirical Strategy

5.1 Overview

Previous work, such as Benkard (2004), has shown that maximum-likelihood approaches to estimating the parameters of dynamic models can be computational demanding, due to the necessity of having to solve for an equilibrium at every guess of the parameter vector. Furthermore, the presence of multiple equilibria requires the econometrician to both compute the set of all possible equilibria and to specify how agents decide on which equilibrium will be played in the data, as in Bajari, Hong, and Ryan (2010).²⁰

In order to sidestep these two issues, I follow the two-step empirical strategy laid out in Bajari, Benkard, and Levin (2007), hereafter referred to as BBL. The intuition of BBL is straightforward: the econometrician lets the agents in the model solve the dynamic program, and finds parameters of the underlying model such that their behavior is optimal. The BBL estimator proceeds in two steps. In the first step, the econometrician flexibly estimates the equilibrium policy functions, $\sigma(s)$. Without imposing any structure, this step simply characterizes what firms do mechanically as

²⁰Borkovsky, Doraszelski, and Kryukov (2008) outline a general approach to solving for the equilibria of Markovian games, and provide a good discussion of why it is generically hard to find all of the equilibria to these systems.

a function of the state vector; these are reduced-form regressions correlating actions to states. This step also avoids the need to compute the equilibrium to the model, as the policy functions are estimated from the equilibrium that is actually played in the data.

The second step is to impose optimality on these recovered policy functions by appealing to the definition of MPNE in Equation 9. By the construction of the value function in Equation 7, given an estimate for $\sigma(s)$ it is possible to construct $EV_i(s; \sigma(s), \theta)$ for some guess of θ . It is possible to construct $EV_i(s; \tilde{\sigma}_i, \sigma_{-i}(s), \theta)$ in analogous fashion by using an alternative policy function for firm i . Since the MPNE requirement holds for all possible alternative strategies, the alternative strategy can be any perturbation of the policies observed in the data, which are held to be optimal under the assumption of profit-maximizing behavior. Given a sufficiently rich set of alternative policies, the BBL estimator finds parameters θ such that profitable deviations from the optimal policies are minimized.

In this application, the first step is to recover the policy functions governing entry, exit, and investment along with the product market profit function.²¹ In the second step, I take these functions and impose the restrictions of the MPNE to recover the dynamic parameters governing the costs of capacity adjustment and exit. Taken collectively, these estimates then allow me to simulate the value of a new firm entering the market, which can be used to recover the distribution of the sunk costs of entry.

The approach in BBL has several regularity assumptions in order to produce valid estimates of the model primitives. Aside from functional form assumptions made below, the following assumption will allow me to group together all markets when estimating policy functions in the first step:

Assumption 1. *The same equilibrium is played in all markets.*

This assumption is critical to obtaining consistent estimates of the unknown parameters. For example, suppose that two equilibria are played in the data, with each equilibrium producing a distinct set of policy functions: $\sigma_1(s)$ and $\sigma_2(s)$. By grouping all markets together, the resulting estimate for $\sigma(s)$ is then a convolution of the policy functions corresponding to the two equilibria, and consistent with neither. It

²¹The assumption that demand and production are static implies that their corresponding parameters can be estimated directly from the data. This is useful because it improves the statistical efficiency of the estimator, as a subset of the parameters are identified independently of those which depend on the construction of the continuation value.

follows directly that the imposition of the MPNE requirement under an inconsistent estimate of $\sigma(s)$ will generically not produce consistent estimates of the underlying primitives. Under Assumption 1, it is possible to estimate the policy functions by grouping data from all markets.²²

I also make the following assumption regarding the beliefs of the firms with respect to the change in regulatory policy.

Assumption 2. *Firms assume that the regulatory environment is permanent.*

This assumption allows me to avoid having to model the beliefs of the firms regarding the distribution of future regulatory environments. In principle, it is possible to model these beliefs if there is an observable covariate which moves around beliefs of possible regulatory changes to the economic environment in the future. However, I assume that the firms behave as if the cost changes due to the Amendments were unanticipated, one-time changes that will never be repeated in the future. This assumption has been used in other applications with regulatory change, such as Rothwell and Rust’s (1995) study of nuclear power plant responses to regulatory change following the Three Mile Island accident.

5.2 Step One: Product Market Profits and Policy Functions

In the first step, I estimate the profits accruing to firms in each period and characterize the entry, exit, and investment behavior of firms conditional on the state variables.

5.2.1 Cement Demand

I estimate several variations on the following specification of the demand for cement in market j at time t :

$$\ln Q_{jt} = \alpha_0 + \alpha_1 \ln P_{jt} + \alpha_{2j} + \alpha'_{3t} X_{jt} + \epsilon_{jt}. \quad (10)$$

The coefficient on market price, α_1 , is the elasticity of demand, and X is a vector of covariates that influence demand. I assume that shocks to demand are iid. I

²²This is a stronger condition than the two-step approach requires: one could estimate the policy functions separately on each market, and then impose the MPNE conditions. However, the limitations of my data preclude such an approach.

instrument for the potential endogeneity of price with the error term using supply-side cost shifters: coal prices, gas prices, electricity rates, and wage rates. Each market has a demand shifter in the intercept, α_{2j} . I estimate several specifications of the demand function, including controls for housing permits, time trends, and population.

5.2.2 Production Costs

In order to estimate the costs of production, I search over δ to match the observed quantities for each firm in each market. For each guess of δ , I solve for the vector of capacity-constrained Cournot quantities. To obtain an interior solution where all firms produce positive quantities, I make the following assumption:

Assumption 3. *The fixed costs of operation, δ_0 , are zero.*

This assumption is empirically driven. As I do not directly observe profits, the only way to infer the fixed costs of operation is to observe firms shutting down production in some periods. In the sample period, all firms produce in all periods, so I cannot identify fixed costs of operation. This normalization does not seem too stringent in the cement industry, both by the economic reasoning that fixed costs cannot be too large as zero production is never observed, and on the a priori grounds that these plants have relatively small staffing requirements and have a production technology where output quantity is directly proportional to energy and material inputs. The production game has an easily-computed fixed point, as the best-response curves are downward sloping in rivals' production.

For each firm i in market by j at time t , the estimator minimizes the sum of squared differences between the observed quantities and the predictions of the model. There are three basic parameters in the cost function: δ_1 , δ_2 , and ν . I also include post-1990 dummy shifters on each of those parameters to capture any changes in the production costs arising from the passage of the Amendments. In order to restrict the threshold at which capacity costs bind to be between 0 and 1, I make use of a logit transformation, $\nu = \exp(\tilde{\nu}) / (1.0 + \exp(\tilde{\nu}))$, and estimate $\tilde{\nu}$. If a firm has multiple plants in a single market, I treat that firm as having a single plant with capacity equal to the sum of capacity in each of those facilities.

5.2.3 Investment Policy Function

Both the presence of fixed costs in the model and empirical evidence suggest that the empirical policy function should be flexible enough to account for lumpy investment behavior. One model that satisfies both of these requirements is the (S, s) rule of investment, such as in Scarf (1959), where firms tolerate deviations from their optimal level of capacity due to fixed adjustment costs.²³ Under the (S, s) rule, firms have a target level bounded on either side by an adjustment band. When the actual level of capacity hits one of the bands, the firm will make an adjustment to the target level.²⁴

I follow Attanasio’s (2000) empirical model of the (S, s) rule, and focus on the investment behavior of firms with positive capacity levels at the start and end of each period. Firms have a target level of capacity, s_{it}^* , which they adjust to when they make an investment:

$$\ln s_{it}^* = \lambda_1' bs(s_{it}) + \lambda_2' bs \left(\sum_{j \neq i} s_{jt} \right) + u_{it}^* \quad (11)$$

where the desired (logged) level of capacity is a function of the firm’s own capacity, the sum of its competitors capacities, and a mean zero error term, u_{it}^* . Since it is desirable to be as flexible as possible in modeling a firm’s behavior as a function of the state, I use the method of linear sieves to estimate the target equation. In this particular case, the basis functions are cubic b-splines, which are finite-dimensional piecewise polynomials, denoted here by $bs(\cdot)$.²⁵

The critical aspect of the (S, s) rule that generates lumpy investment behavior is that firms only adjust s_{it} to s_{it}^* when current capacity exceeds one of the bands

²³Deriving this rule as the explicit solution to an optimization problem is involved—see Hall and Rust (2000) for an example of the optimality of this rule in an inventory setting. Two observations support the use of the (S, s) rule: the estimated bands have statistically significant gaps from the target level, and when I solve for an equilibrium of the model firms engage in lumpy investment behavior that mirrors the (S, s) rule.

²⁴This model also nests the model of continuous investment as the bands go to zero, and is thus quite flexible in its ability to capture a range of investment behavior.

²⁵Chen (2006) provides an exhaustive overview of b-splines and other linear sieves. Throughout the paper, I use uniform b-splines with ten knots, where the range of the knots was chosen to bound the empirical data. Further implementation details of the b-splines are available from the author upon request.

around the target level. The lower and upper bands are given by:

$$s_{it} = s_{it}^* - \exp \left(\lambda'_3 b s_1(s_{it}) + \lambda'_4 b s_2 \left(\sum_{j \neq i} q_{jt} \right) + \underline{u}_{it}^b \right), \quad (12)$$

and

$$\bar{s}_{it} = s_{it}^* + \exp \left(\lambda'_5 b s_1(s_{it}) + \lambda'_6 b s_2 \left(\sum_{j \neq i} q_{jt} \right) + \bar{u}_{it}^b \right). \quad (13)$$

The inclusion of the exponential function ensures that the desired level of adjustment is always in between the bands. This model also nests a model of continuous adjustment as the width of the bands goes to zero. I assume that the residuals in the bands are iid normal with zero mean and equal variance, and are independent of the error in the target. When a firm makes an adjustment, it reveals both the target level, s_{it}^* , and the size of the band, $s_{it}^* - s_{i,t-1}$, which is sufficient to identify the parameters of the adjustment policy function.

I assume that the upper and lower bands are symmetric functions of the target capacity ($\lambda_3 = \lambda_5$ and $\lambda_4 = \lambda_6$); the reason is that the upper bound is not precisely estimated if treated separately. Since I assume that the change in capacity simultaneously reveals the size of the band and the target level, I use linear regression to recover λ . I estimate separate policy functions for the period before 1990, and the period after 1990. This will capture any differences in the firms' equilibrium investment behavior caused by a permanent shift in the cost and regulatory environment.

5.2.4 Entry and Exit Policy Functions

I characterize the probability of entry using a probit regression:

$$Pr(\chi_i = 1; s_i = 0, s) = \Phi \left(\psi_1 + \psi_2 \left(\sum_{j \neq i} s_{jt} \right) + \psi_3 1(t > 1990) \right), \quad (14)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal, and $1(\cdot)$ is the indicator function. The exit policy is also modeled analogously:

$$Pr(\chi_i = 0; s_i > 0, s) = \Phi \left(\psi_4 + \psi_5 s_{it} + \psi_6 \left(\sum_{j \neq i} s_{jt} \right) + \psi_7 1(t > 1990) \right). \quad (15)$$

Explanatory variables in both policy functions are a constant, the sum of competitors' capacities, and a dummy variable for before and after 1990.²⁶ I also add the firm's own capacity to the exit equation. I am assuming that there is only one possible entrant in any period; given the very low rate of entry in the cement industry, this assumption is not very important.

5.3 Step Two: Recovering the Structural Parameters

The first step provides functions that describe both how the state vector evolves over time and what product market profits are at each state. The second step finds parameters that make these observed policy functions optimal, given the underlying theoretical model.²⁷

The per-period payoff function for a firm is:

$$\begin{aligned} \pi_i(s, \sigma(s); \theta, \epsilon_i) = & \bar{\pi}_i(s) \cdot 1 - 1(x_i > 0) \cdot (\gamma_{1i} + \gamma_2 x_i + \gamma_3 x_i^2) \\ & + 1(x_i < 0) \cdot (\gamma_{4i} + \gamma_5 x_i + \gamma_6 x_i^2) + 1(exit) \cdot \phi_i, \end{aligned} \quad (16)$$

where the indicator functions represent which of the discrete actions the firm has undertaken in that period. Integrating out the private shocks results in the following per-period payoff function:

$$\begin{aligned} E_{\epsilon_i} \pi_i(s, \sigma(s); \theta) = & \bar{\pi}_i(s) \cdot 1 - p_i(s) \cdot (\tilde{\gamma}_{1i} + \gamma_2 x_i + \gamma_3 x_i^2) \\ & + p_d(s) \cdot (\tilde{\gamma}_{4i} + \gamma_5 x_i + \gamma_6 x_i^2) + p_e(s) \cdot \tilde{\phi}_i, \end{aligned} \quad (17)$$

where indicator functions for investment, divestment, and exit have been replaced with their associated equilibrium probabilities, $p_i(s)$, $p_d(s)$, and $p_e(s)$, respectively. The draws of private information have been replaced by their conditional expectations. The tilde emphasizes that the expected values of these draws are not equal to their unconditional means of their underlying distributions, as firms only undertake actions when the associated shock is sufficiently favorable.

The conditional mean of exit costs is the simplest case. Recalling that the firm

²⁶Ideally, one would recover these probabilities using a very flexible function of the state variables. However, exploration of more flexible functional forms did not lead to better statistical fits, likely due to the relatively limited amount of variation in the data.

²⁷This section follows the derivations in Bajari, Benkard, and Levin (2007) closely.

does not know the draws of investment and divestment fixed costs when it makes its choice, the probability that a firm exits is:

$$Pr(exit|s) = Pr(\phi_i > E \max [V_i^+(s) - \gamma_{1i}, V_i^-(s) - \gamma_{4i}, V_i^0(s)]), \quad (18)$$

where $V_i^+(s)$, $V_i^-(s)$, and $V_i^0(s)$ are the values associated with (optimal) investment, divestment, and doing nothing, respectively. Since the probability of exit encapsulates all of the relevant information facing the firm at a specific state, it follows that the conditional mean of exit costs is also solely a function of the probability of exit:

$$\tilde{\phi}_i = E[\phi_i | \phi > E \max [V_i^+(s) - \gamma_{1i}, V_i^-(s) - \gamma_{4i}, V_i^0(s)]] = \theta_\phi \cdot bs(p_e(s)), \quad (19)$$

where I have replaced the unknown conditional mean function with a linear b-spline. The b-spline allows for a flexible approximation to the relationship between the expected value of the truncated distribution and the probability of exit.

The conditional mean of fixed costs of investment, $\tilde{\gamma}_{1i}$, is slightly complicated by presence of a second shock.²⁸ The probability a firm invests is:

$$Pr(invest|s) = Pr(V_i^+(s) - \gamma_{1i} \geq V_i^-(s) - \gamma_{4i}, V_i^+(s) - \gamma_{1i} \geq V_i^0(s)). \quad (20)$$

This probability depends on both value functions and the draw of the divestment fixed cost. Therefore, in principle the conditional mean is also a function of these two probabilities:

$$\tilde{\gamma}_{1i} = E[\gamma_{1i} | V_i^+(s) - \gamma_{1i} \geq V_i^-(s) - \gamma_{4i}, V_i^+(s) - \gamma_{1i} \geq V_i^0(s)] = \theta_{\gamma,1} \cdot bs(p_i(s), p_d(s)), \quad (21)$$

where $bs(p_i(s), p_d(s))$ is a tensor product of linear b-splines defined over the unit square. The value of γ_{1i} that the firm draws from is truncated above by the minimum of its competing alternatives. Intuitively, as the other alternatives become more attractive, as reflected in an increasing probability of choosing those alternatives, the draw of investment costs that would induce a firm to undertake that action has to become more favorable.

In practice, I estimate Equation 21 (and the associated conditional mean function for divestment costs) using functions of only the associated action's choice probability.

²⁸The conditional mean for divestment is symmetric.

The reason is data-driven: under the (S, s) rule estimated above, the probability of seeing any state with positive probability on both investment and divestment is so small that the computer is incapable of differentiating the probability from zero. This implies that the conditional probability of investment or divestment is equal to the unconditional probability.²⁹

Following Bajari, Benkard, and Levin (2007), I leverage the fact that all of the unknown parameters enter linearly into the payoffs of the firm. Equation 17 can be rewritten as the inner product of a row vector and a column vector:

$$\pi_i(s, a; \theta) = [\bar{\pi}_i(s_{it}) \quad \zeta(s_{it})] \cdot [1 \quad \theta]'. \quad (22)$$

The per-period payoff function is completely known, and therefore is only multiplied by one. Defining the following:

$$W_i(s_t; \sigma(s)) = E_{\sigma(s)} \sum_{t'=0}^{\infty} \beta^{t'} [\bar{\pi}_i(s_{i,t+t'}) \quad \zeta(s_{i,t+t'})], \quad (23)$$

the value function is then:

$$V_i(s_t; \sigma(s), \theta) = W_i(s_t; \sigma(s)) \cdot [1 \quad \theta]'. \quad (24)$$

Imposing the Markov perfect equilibrium condition (see Equation 9) for all alternative policies $\tilde{\sigma}_i$ obtains:

$$W_i(s_o; \sigma_i^*, \sigma_{-i}) \cdot [1 \quad \theta]' \geq W_i(s_o; \tilde{\sigma}_i, \sigma_{-i}) \cdot [1 \quad \theta]', \quad (25)$$

At the true parameters the above relation should hold for all alternative policies. Exploiting the linearity of the unknown parameters, I can rewrite the above equation in terms of profitable deviations from the optimal policy:

$$g(\tilde{\sigma}_i; \theta) = [W_i(s; \tilde{\sigma}_i, \sigma_{-i}) - W_i(s; \sigma_i^*, \sigma_{-i})] \cdot [1 \quad \theta]'. \quad (26)$$

To implement the estimator, I draw $n_k = 1,250$ alternative policies to generate a set of inequalities by adding noise to the optimal policy functions. For example, to perturb the exit policy function I add an error drawn from the standard normal to

²⁹I show that this is sufficient to recover the distribution of fixed costs in Appendix A.

the terms inside the exit probit. The estimator then searches for parameters such that profitable deviations from the optimal policies are minimized:

$$\min_{\theta} Q_n(\theta) = \frac{1}{n_k} \sum_{j=1}^{n_k} 1(g(\tilde{\sigma}_{i,j}; \theta) > 0) g(\tilde{\sigma}_{i,j}; \theta)^2. \quad (27)$$

The linearity of the unknown parameters becomes useful during the minimization, as I do not have to recompute separate outcome paths for each set of parameters. The function is not trivially minimized at the zero vector because the profits from the product market enter in each time period.³⁰ I use the Laplace-type estimator (LTE) of Chernozhukov and Hong (2003) to search over θ in Equation 27.

5.3.1 Distribution of Sunk Entry Costs

Having recovered the policy functions and the parameters necessary for the construction of the period payoffs it is possible to find the distribution of sunk costs. Consider the value function of a potential entrant:

$$V_i^e(s; \sigma(s), \theta, \epsilon_i) = \max \left\{ 0, \max_{x_i \geq 0} [-\kappa_i - \gamma_{1i} - \gamma_2 x_i - \gamma_3 x_i^2 + \beta E(V(s')|s, \sigma(s))] \right\}. \quad (28)$$

All of the terms in Equation 28 are known or computable except for the distribution of κ_i . Assuming the entry costs are distributed normally with mean μ_κ and variance σ_κ^2 , the probability that a firm enters is equal to the probability of that entrant receiving a draw of the sum of two fixed costs that is less than the value of entry:

$$Pr(\kappa_i + \gamma_{1i} \leq EV^e(s)) = \Phi(EV^e(s); \mu_\kappa + \mu_\gamma^+, \sigma_\kappa^2 + \sigma_\gamma^{+2}), \quad (29)$$

where Φ is the normal CDF.³¹ The left-hand side of Equation 29 corresponds to the entry policy function estimated in Equation 14, while $EV^e(s)$ can be computed through forward simulation. Drawing $s = \{1, \dots, NS\}$ random states of the industry,

³⁰For example, all perturbed policy functions that lead to higher investment would increase profits by gaining the firm a larger market share but would not incur any investment costs at the zero vector; this implies that firms would invest to arbitrarily large capacity, which is inconsistent with their observed equilibrium behavior.

³¹The right-hand side follows from the fact that the distribution of a sum of two normally-distributed variables is also normal with mean and variance equal to the sum of the addends' means and variances.

Table 3: Cement Demand Estimates

	I	II	III	IV	V	VI
Price	-3.21 (0.361)	-1.99 (0.285)	-2.96 (0.378)	-0.294 (0.176)	-2.26 (0.393)	-0.146 (0.127)
Intercept	21.3 (1.52)	10.30 (1.51)	20.38 (1.56)	-3.41 (1.09)	11.6 (2.04)	-6.43 (0.741)
Log Population		0.368 (0.0347)		0.840 (0.036)	0.213 (0.074)	0.789 (0.033)
Log Permits					0.218 (0.072)	0.332 (0.035)
Market Fixed Effects	No	No	Yes	Yes	No	Yes

Dependent variable is logged quantity. Instruments were gas prices, coal prices, electricity prices, and skilled labor wage rates. There are a total of 517 observations.

I search for parameters of this distribution which match the observed probabilities of entry as well as possible:

$$\min_{\{\mu_\kappa, \sigma_\kappa^2\}} \frac{1}{NS} \sum_{i=1}^{NS} [Pr(\text{entry}|s_i) - \Phi(EV^e(s_i); \mu_\kappa + \mu_\gamma^+, \sigma_\kappa^2 + \sigma_\gamma^{+2})]^2. \quad (30)$$

I estimate the parameters of the distribution of sunk entry costs separately for the time periods before and after the 1990 Amendments.

5.3.2 Standard Errors

Standard errors were calculated by random subsampling without replacement at the market-history level, as in Politis and Romano (1994). I randomly drew subsamples of 19 complete market histories 500 times.

6 Empirical Results

Demand Curve I estimate the parameters in Equation 1, the demand curve for Portland cement, using market-level data on prices and quantities. I use several cost-side shifters serving as instruments to account for the endogeneity of prices. The results are presented in Table 3.

The first specification is the simplest, as it has no covariates. The price elasticity of demand is precisely estimated to be -3.21. However, one may expect that

demand may vary across markets due to population or other unobservable factors. The next two models test for these factors. Specification II adds in controls for population, in logs, which is estimated to have a positive effect on quantity. The elasticity falls to -1.99, and the intercept is substantially lower. The average log population in the sample is 15.87; multiplying by the coefficient on log population shifts the constant back to 16.1, closer to the baseline model's intercept. Specification III includes market-specific fixed effects in lieu of population shifters. The results are very similar to the baseline specification, with higher elasticity and intercept than in specification II. Specification IV includes both population and fixed effects. The result is a statistically insignificant estimate of the price coefficient, a negative intercept, and a positive coefficient on log population. Specification V includes a measure of housing permits allocated in each market-year. The elasticity of demand is estimated to be -2.26, with positive coefficients on population and permits. Specification VI includes fixed effects for markets as well, which leads to a small and statistically insignificant estimate of the price elasticity.

The specification I choose to use throughout the rest of the analysis is specification III. The reasoning behind this choice is three-fold. First, it appears that market fixed effects capture much of the same cross-market variation in prices that population and permits do. A regression of quantities on prices, population, permits, and the interaction of population and permits with a time trend leads to imprecisely measured zeros on the interaction terms. This suggests that population and permits are not changing very much within market, and their explanatory power is cross-sectional. This argues that market fixed effects may reasonably proxy for these effects. Second, while the fixed effects approach is not as nuanced as the population and permits approach, which utilizes more data, it has the benefit of being the more parsimonious specification. With limited data, as in the present context, this is a strength. Finally, a simple plausibility check suggests that the specification with the higher elasticity results in a more reasonable estimate of plant costs. If one takes a specification with a lower elasticity as the demand curve, and works through the ensuing empirical exercise, the resulting estimates imply that firms face unreasonably large investment costs in order to rationalize their behavior. Otherwise, firms would be leaving very large amounts of money on the table; as such, the estimator predicts investment costs on the order of several billion dollars for a modern plant, which is inconsistent with

anecdotal newspaper evidence and the accounting data cited in Salvo (2010).³² For these reasons, I proceed with specification III as the model for cement demand.

To verify that the instruments used in the demand estimation are both correlated with the endogenous regressor and orthogonal to the error term, I evaluate both the fits of the instruments on the endogenous regressor and the Anderson-Rubin statistic. The F-statistic of the first-stage regression of the instruments on the endogenous regressor results is 42.99, which is significant at the one percent level and well above the rule-of-thumb threshold of 10. The Anderson-Rubin statistic is 52.54, which is also significant at the one percent level. I conclude that the tests fail to reject the hypothesis that the instruments were both well-correlated with prices and orthogonal to the error terms.

Finally, I test for the presence of time trends in each of the markets. While the F-test rejects the hypothesis that all of the coefficients are significantly different from zero, most of the market-time trends are not individually significant (22 out of 26 markets). The elasticity of demand is precisely estimated to be -2.26, with an intercept of 17.51. Saturating the model with trends and dummy variables is strong empirical evidence that the elasticity is in the range of -2 to -3, as the explanatory variables account for a wide range of market- and time-specific unobserved heterogeneity.³³

Production Costs Having estimated the demand curve, I recover the production cost parameters by matching predicted quantities as closely as possible to their empirical counterparts. I estimate six parameters: marginal cost, capacity cost, the capacity binding level, and post-1990 shifters for each. The results are shown in Table 4. The estimates indicate that capacity costs become important as firms increase production beyond 87 percent of their boilerplate capacity. Once firms cross this threshold they experience large, linearly increasing marginal costs as they cut into the normal period of maintenance downtime. The penalty for cutting out your maintenance is significant, preventing most producers from exceeding 90 percent of their stated production capacity. The shifters for post-1990 are not statistically significantly different from zero at the 5 percent level. This bolsters the argument that the Amendments did not have an influence on marginal costs.

As a check on the estimated parameters, I compute the market price, revenues,

³²A table of announced plant costs is available in a previous version of the paper and from the author upon request.

³³Additional specifications are available from the author upon request.

Table 4: Production Function Results

<u>Production Function Estimates</u>		
Parameter	Coefficient	Standard Error
Marginal Cost (δ_1)	31.58	1.91
Capacity Cost (δ_2)	1.239	0.455
Capacity Cost Threshold ($\tilde{\nu}$)	1.916	0.010
Marginal Cost post-1990 shifter	2.41	3.33
Capacity Cost post-1990 shifter	-0.0299	0.22
Capacity Cost Threshold post-1990 shifter	0.0917	0.0801
<u>Prices, Revenues, and Profits</u>		
Variable	Value	Standard Deviation
Price	57.81	16.83
Revenues	39,040	19,523
Costs	22,525	11,051
Profit	16,515	12,244
Margin	39.29 percent	18.21 percent

The binding threshold at which the capacity costs become important is restricted to $[0, 1]$ by estimating a logit probability: $\nu = \exp(\tilde{\nu}) / (1.0 + \exp(\tilde{\nu}))$. At the estimated value of 1.916, this implies that capacity costs start to bind at an approximately 87 percent utilization rate.

costs, and profit margin for every firm in my sample. The summary statistics for these values are shown in Table 4. The prices are well within the range seen in the data. The average firm grosses slightly less than \$40M a year. Profits average just over \$16M a year, which is little less than a 40 percent profit margin. This is a plausible gross return, as public financial records for major cement producer Lafarge North America report an 33 percent average gross profit margin for the three-year period 2002-2004.³⁴

To test the assumption that the firms have no persistent productivity differences, I regressed output quantity on own capacity and various controls. If there are productivity differences across firms, it should be expressed in their ability to utilize their productive capacity: more productive firms produce more given the same amount of capacity. This should be an especially strong test given that most firms were capacity constrained during this time period. The controls include the capacity of rival firms, a time trend, market fixed effects, and capacity interacted with dummy variables for whether the firm entered or exited during the sample period. Market fixed effects capture variation in local demand conditions, and the capacity of rival firms shifts around residual demand facing any firms. The dummy variables for entry and exit

³⁴Sales and profit data are from Hoover's Online "Annual Financials" fact sheet for Lafarge S.A., 2002-2004. <http://www.hoovers.com>.

Table 5: Productivity Estimates

Specification	I	II	III	IV	V
Capacity	0.8617 (0.002)	0.8600 (0.002)	0.860 (0.002)	0.860 (0.002)	0.860 (0.002)
Rivals' Capacity	-0.007 (0.001)	-0.005 (0.001)	-0.002 (0.001)	-0.003 (0.001)	0.0003 (0.0006)
Firm Entered * Capacity		0.0009 (0.0027)	0.0002 (0.0027)	0.0112 (0.0064)	0.0103 (0.007)
Firm Exited * Capacity		-0.0154 (0.0035)	-0.0128 (0.0036)	-0.0173 (0.0078)	-0.0135 (0.008)
Time Trend			0.671 (0.130)	0.681 (0.131)	
Entry Dummy				-11.66 (6.141)	-11.49 (6.678)
Exit Dummy				3.041 (4.810)	0.492 (5.107)
Market Fixed Effects	Yes	Yes	Yes	Yes	No
Market-Time Fixed Effects	No	No	No	No	Yes
R^2	0.9925	0.9925	0.9926	0.9926	0.9933

Number of observations = 2,233.

capture systematic differences in productivity, as measured by production per unit of capacity. The results are presented in Table 5.

The first specification considers only own capacity and rivals' total capacity, controlling for market with fixed effects. This model does remarkably well in fitting the data, where $R^2 = 0.9925$. This suggests that unobservable productivity differences cannot be very important in the sample, as otherwise there would be significant variance in the output that variation in capacity alone could not explain. This variable is very precisely estimated across all the specifications, further supporting the idea that production is largely explained by observable capacity, controlling for factors common to all firms in the market.

To examine the question of whether firms were selecting in and out of the cement industry along these unobservables, I included dummy variables for firms that entered and exited during the sample period. The second specification allows the production rate per unit of capacity to shift for firms that entered and exited. The results suggest that new firms are no more productive per unit capacity than the average firm in the industry, while exiting firms on average produced 1.8 percent less per unit of capacity than the average firm. The third specification adds a time trend to production; if firms tended to exit earlier in the sample vis a vis new entrants, this could bias the selection effect. The addition of a time trend reduces the difference in productivity for exiting firms, although it is still significant. The last specification also adds the

dummy variable for entering and exiting firms directly. If the productivity differences are explained by differences in startup times, this could show up through a level effect. Curiously the results suggest that entering firms have lower production levels than both the average firm and exiting firms, which are more productive than average firms. The production per unit capacity differential for entering (exiting) firms is still positive (negative), although neither is now significant at the 5 percent level. Finally, saturating the model with time-market fixed effects to control for any other observable factors leads to insignificant estimates for productivity differences across entering and exiting firms. On the whole, it appears that there is little reduced-form evidence for productivity differences across entering firms and the average firm in the sample.

Investment Policy Estimates for several specifications of the (S, s) rule are presented in Tables 6 and 7.

The first two specifications in the band equation use levels of capacity as explanatory variables, with specification II including population as a level shifter. I constructed basis functions of all three variables to allow the marginal effect to vary with the magnitude of the covariate. The base specification does a very good job of matching the observed size of adjustments. The adjusted R^2 for both regressions is almost 0.9. A regression of fitted values on the actual adjustments produces a regression with an imprecisely estimated zero intercept and a tightly estimated slope parameter of almost exactly 1; this indicates that on average the model is able to fit the observed gaps very well using the flexible basis functions over the sum of competitors' capacity and a firm's own capacity. The inclusion of population basis functions, as in specification II, slightly improves the fit of the model at the expense of greatly increasing the variance of the estimated parameters. A standard F-test fails to reject the null hypothesis that the coefficients on the population basis functions are jointly zero (p-value of 0.2149).

One may think that the inclusion of population should be through per-capita capacities to adjust for differences in market sizes. Specification III runs the same regression of log adjustment on basis functions of per-capita capacities. In this case the model does slightly less well in fitting the size of the adjustment band. The adjusted R^2 is lower, and the estimated variance in the error term is a bit larger. The signs on the coefficients for own capacity are reversed, but are not statistically

Table 6: Investment Policy Function Results: Adjustment Band Size

Specification	I	II	III	IV
Sum Competitors Capacity B-spline 1	6.31 (0.973)	7.18 (5.42)	3.29 (0.827)	5.74 (1.14)
Sum Competitors Capacity B-spline 2	6.51 (0.930)	7.16 (5.43)	2.04 (0.953)	4.53 (1.37)
Sum Competitors Capacity B-spline 3	5.66 (0.910)	6.41 (5.40)	3.57 (0.888)	4.89 (1.28)
Sum Competitors Capacity B-spline 4	6.98 (0.960)	7.85 (5.46)	2.05 (0.978)	4.05 (1.36)
Sum Competitors Capacity B-spline 5	5.77 (0.939)	6.72 (5.33)	2.91 (0.994)	5.40 (1.47)
Sum Competitors Capacity B-spline 6	7.3 (0.944)	8.15 (5.67)	2.11 (1.15)	4.23 (1.50)
Own Capacity B-spline 1	-3.79 (0.923)	-3.97 (0.925)	0.374 (0.880)	-2.71 (1.22)
Own Capacity B-spline 2	-3.3 (0.893)	-3.37 (0.894)	0.720 (0.902)	-0.754 (1.22)
Own Capacity B-spline 3	-2.3 (0.967)	-2.51 (0.969)	1.06 (1.04)	-0.325 (1.28)
Own Capacity B-spline 4	-1.72 (0.943)	-1.76 (0.952)	1.87 (1.27)	-0.149 (1.60)
Own Capacity B-spline 5	-2.63 (1.35)	-2.66 (1.35)	2.05 (2.25)	3.32 (2.02)
Population B-Spline 1		-5.11 (6.78)		
Population B-Spline 2		0.886 (5.16)		
Population B-Spline 3		-1.39 (5.52)		
Population B-Spline 4		-0.008 (5.06)		
Population B-Spline 5		-1.60 (6.69)		
Capacity is Per-Capita	No	No	Yes	Yes
Region Fixed Effects	No	No	No	Yes
Adjusted R^2	0.8952	0.8955	0.8816	0.8946
Band σ^2	1.40	1.40	1.56	1.41

Dependent variable is the natural log of the change in capacity. Number of capacity changes = 774. Parameters estimated using OLS. Capacity is measured in thousands of tons per year. Population is denominated in tens of millions.

significant. Competitors' aggregate capacity also enters in strongly and positively. To ensure that this isn't proxying for market-level demand shifters that make large firms uniformly more attractive, Specification IV adds in region fixed effects. The fixed effects are estimated very imprecisely, and tend to be small deviations around zero. The magnitude of the coefficients on aggregate competitor capacity are even stronger with the fixed effects, and the sign of own capacity reverts to being negative, although they are estimated with significant amount of variance. While this specification fits better than the per-capita model without fixed effects, it still is not as good a fit as specification I, which hereafter is the preferred empirical specification.

Table 7 reports four specifications of the equation for the level a firm desires to adjust to given it is going to make an investment. The fit in all the specifications is extremely tight; the lowest adjusted R^2 is 0.9958. The estimated variance of the error term in all specifications is also very low. These statistics suggest that the model is capable of accurately fitting the capacity levels that firms adjust to quite tightly as a flexible function of competitors' aggregate capacity and a firm's own capacity. As in the band equation, I test several specifications. In the baseline specification, the explanatory variables are b-splines of aggregate competitor capacity and own capacity. The results suggest that target values are strictly increasing across the range of competitor capacities seen in the data. The function is negative with respect to firm's own capacity, although the coefficients are decreasing in the spline, which suggests that larger firms prefer to make larger adjustments.

This result may be due to the fact that larger firms operate in larger markets, and therefore the residual demand curve is larger. To test this hypothesis, specification II includes b-splines of market population. The coefficients on these variables are positive, which supports the idea that larger markets support higher investment. However, the coefficients on own capacity are virtually exactly the same, suggesting that the pattern of larger firms having larger target levels of capacity holds when controlling for market size. The coefficients on the aggregate competitor capacity decrease by roughly the size of the population variables. Specification III includes region fixed effects, to test for market-level heterogeneity not captured otherwise. The fixed effects are very imprecisely estimated. The overall effect is to push the coefficients on aggregate competitor capacity back toward the their levels in specification I while weakening the effect of own capacity. The population effects are rendered statistically insignificant, implying that market-specific variation in target levels is

Table 7: Investment Policy Function Results: Investment Target

	I	II	III	IV
Sum Competitors Capacity B-spline 1	7.74 (0.124)	5.80 (0.714)	7.26 (0.927)	7.094 (0.282)
Sum Competitors Capacity B-spline 2	7.70 (0.123)	5.67 (0.715)	7.06 (0.929)	6.96 (0.326)
Sum Competitors Capacity B-spline 3	7.76 (0.120)	5.80 (0.711)	7.17 (0.936)	7.50 (0.303)
Sum Competitors Capacity B-spline 4	7.64 (0.127)	5.65 (0.719)	6.60 (0.964)	6.50 (0.334)
Sum Competitors Capacity B-spline 5	7.88 (0.124)	5.96 (0.701)	6.82 (0.987)	7.31 (0.340)
Sum Competitors Capacity B-spline 6	7.59 (0.124)	5.52 (0.746)	6.36 (0.992)	6.71 (0.391)
Own Capacity B-spline 1	-2.24 (0.121)	-2.24 (0.121)	-2.15 (0.124)	-0.912 (0.301)
Own Capacity B-spline 2	-1.36 (0.118)	-1.36 (0.118)	-1.31 (0.124)	-0.136 (0.308)
Own Capacity B-spline 3	-0.752 (0.128)	-0.753 (0.128)	-0.702 (0.130)	-0.762 (0.354)
Own Capacity B-spline 4	-0.182 (0.124)	-0.186 (0.125)	-0.120 (0.134)	1.27 (0.432)
Own Capacity B-spline 5	0.074 (0.179)	0.0096 (0.178)	0.063 (0.181)	-0.831 (0.767)
Population B-Spline 1		1.43 (0.892)	0.482 (2.30)	
Population B-Spline 2		2.08 (0.679)	0.483 (0.876)	
Population B-Spline 3		1.95 (0.727)	0.656 (1.04)	
Population B-Spline 4		1.98 (0.666)	0.015 (0.802)	
Population B-Spline 5		2.37 (0.881)	-0.566 (1.21)	
Capacity is Per-Capita	No	No	No	Yes
Region Fixed Effects	No	No	Yes	No
Adjusted R^2	0.9994	0.9994	0.9995	0.9958
Estimated σ^2	0.0244	0.0242	0.0235	0.184

Dependent variable is log of capacity level after adjustment. Number of capacity changes = 774. Parameters estimated using OLS. Capacity is measured in thousands of tons per year. Population is denominated in tens of millions.

largely captures by the fixed effects.

To examine the effect of population further, specification IV runs the regression with per-capita capacities instead of levels. The average population is about 1.14 when normalized to be in tens of millions of people, which helps explain why the coefficients on the capacity variables are very close to the previous specifications. The pattern of coefficients looks very similar to the first three specifications, while not fitting the data as well; the regression error is an order of magnitude larger.

Since the first three specifications all provide an excellent fit to the data, and specification II and III have a number of statistically insignificant coefficients, I hereafter proceed using specification I as the empirical model for the target level of adjustment.

Entry and Exit Policy Several estimated specifications of the entry and exit policy function results are presented in Table 8. I estimated both policy functions both in absolute levels and in per-capita levels, to control for unobserved market-level variation in demand that could change the policy functions. Specifications I and II estimate the exit policy in levels, with and without controls for population. The sign and magnitude of the estimated coefficients are very close in both specifications. Own capacity decreases the probability of exit, and an increase competitors' capacity increases that probability. Both explanatory variables proxy for the level of residual demand facing the firm when it makes an exit decision. The constant is near -1 in both regressions and the dummy variable for post-1990 is -0.595 and -0.607, respectively. To put these numbers in context in specification I, moving a firm into the post-1990 environment has the same effect as increasing its own capacity by 368 thousand tons per year. The variable with the most explanatory power is the post-1990 dummy variable; the marginal effect in specification I, evaluated at the means of the other explanatory variables, is to decrease the probability of a given firm's exit from 2.1 percent to 4 tenths of a percent, a fivefold decrease. This directly reflects the overall exit rates in the industry: there were 51 exits in the period before 1990 and six exits after, a difference of 81 percent.

The effect of competitor's capacity is estimated to be positive and small under both specifications; the marginal effect is so small as to be economically unimportant. The addition of population as a control to the exit equation improves the fit of the model only marginally; the likelihood ratio test fails to reject the null hypothesis that the coefficient on population is equal to zero, as it is estimated very imprecisely. As

Table 8: Entry and Exit Policy Results

Specification	I	II	III	IV
Exit Policy				
Own Capacity	-0.0015661 (0.000268)	-0.0015795 (0.0002712)		
Competitors Capacity	0.0000456 (0.0000173)	0.0000379 (0.0000249)		
Population		0.0590591 (0.1371835)		
After 1990	-0.5952687 (0.1616594)	-0.606719 (0.1639955)	-0.6328867 (0.157673)	-0.4623664 (0.1910193)
Own Capacity per Capita			-0.0005645 (0.0001255)	-0.0010199 (0.0002164)
Competitors Capacity per Capita			0.0000744 (0.00000286)	0.0002379 (0.0001023)
Constant	-1.000619 (0.1712286)	-1.019208 (0.176476)	-1.664808 (0.1475588)	-1.529715 (0.3526938)
Region Fixed Effects	No	No	No	Yes
Log Likelihood	-227.21	-227.12	-238.54	-217.38
Entry Policy				
Competitors Capacity	0.0000448 (0.0000365)	-0.0003727 (0.0002351)		
After 1990	-0.6089773 (0.2639545)	-0.8781589 (0.3229502)	-0.602279 (0.2651052)	-1.003239 (0.337589)
Constant	-1.714599 (0.2152315)	-0.454613 (0.7086509)	-1.665322 (0.2642566)	-0.3434765 (0.6624767)
Competitors Capacity per Capita			0.000026 (0.000038)	-0.0003633 (0.0001766)
Region Fixed Effects	No	Yes	No	Yes
Log Likelihood	-70.01	-56.47	-70.491	-55.53
Prob > χ^2	0.0177	0.4516	0.0287	0.3328

Sample size for exit policy function = 2233; sample size for entry policy function = 414. Capacity is measured in thousands of tons of cement per year. Population is normalized to be measured in tens of millions. Per capita capacity is measured as thousands of tons per year per tens of millions in population.

a result, the marginal effects in specification II are very similar to specification I.

One could be concerned that these specifications fail to adequately capture factors that influence the residual demand curve. To guard against this, I also estimated exit policies that were functions of per-capita capacity, with and without region fixed effects. The results are shown in columns III and IV. The per-capita results roughly mirror the ones above—the post-1990 dummy still dominates the effects of the other two variables. The relative marginal effect of the post-1990 shifter is stronger, own-capacity weaker, and competitor’s capacity stronger in the per-capita model. The inclusion of region fixed effects improves the fit of the model, although the likelihood ratio test fails to reject the hypothesis that the coefficients on the fixed effects are jointly zero. I also estimated a model where the competitors’ per-capita aggregate capacity was expanded to a fourth-degree orthogonal polynomial. The results of this model were almost identical to those from specification III. The inclusion of a time trend did not change the results. The fit of specification III, as measured by the improvement in the likelihood function after including covariates, is not as good than specification I, so I take specification I to be my preferred empirical model.

I also estimated several specifications of the entry policy functions, shown in the bottom panel of Table 8. The baseline rate of entry is low, as accounted for by the constant, which is estimated to be negative in all specifications. The post-1990 dummy is negative and significant in all specifications. Analogous to the exit policy function, this reflects the empirical trends for entry; there were 15 entries in the period before 1990 and four entries after the passage of the Amendments in 1990. The signs on incumbent capacity, whether in levels or per-capita, are positive, small, and statistically insignificant in specifications without region fixed effects. The inclusion of fixed effects flips the sign on incumbent capacity negative, as expected, but a likelihood ratio test fails to reject the hypothesis that all coefficients are equal to zero. I have reported the p-values for this test below each specifications log-likelihood value in the last row. The only model that is not rejected at the two percent level is specification I, which estimates entry as a function of incumbent capacity. The per-capita specification is modestly less significant than specification I; therefore I use the model in levels for the remainder of the analysis.

Dynamic Parameters The results of the second step estimation described above are presented in Table 9. The projection of the b-spline coefficients onto their un-

Table 9: Dynamic Parameters

Parameter	Before 1990		After 1990		Difference	
	Mean	SE	Mean	SE	Mean	SE
Investment Cost	230	85	238	51	-8	19
Investment Cost Squared	0	0	0	0	0	0
Divestment Cost	-123	34	-282	56	-155	35
Divestment Cost Squared	3,932	1,166	5,282	1,130	1,294	591
Investment Fixed Costs						
Mean (μ_{γ}^+)	621	345	1,253	722	653	477
Standard Deviation (σ_{γ}^+)	113	72	234	145	120	97
Divestment Fixed Costs						
Mean (μ_{γ}^-)	297,609	84,155	307,385	62,351	12,665	34,694
Standard Deviation (σ_{γ}^-)	144,303	41,360	142,547	29,036	109	17,494
Scrap Values						
Mean (μ_{ϕ})	-62,554	33,773	-53,344	28,093	9,833	21,788
Standard Deviation (σ_{ϕ})	75,603	26,773	69,778	27,186	-6,054	11,702
Entry Costs						
Mean (μ_{κ})	182,585	36,888	223,326	45,910	43,654	21,243
Standard Deviation (σ_{κ})	101,867	22,845	97,395	14,102	-6,401	12,916

Means of the parameters are reported for the pre-1990 period and the post-1990 period. Units are in thousands of dollars per ton for capital costs; the distributions are denominated in thousands of dollars. Standard errors were calculated via subsampling.

derlying distributions are reported. The fixed costs of investment are significant, reported at \$620,000 in the pre-1990 period and doubling to \$1.25M in the post-1990 period, although this difference is not significant. The fixed costs of adjustment are relatively small next to the variable costs of investment for a typical plant. The early and late estimates for the marginal cost of adjustment are very close, \$230 per ton before 1990 and \$238 per ton after 1990, and statistically indistinguishable. These costs imply that a 1.5M plant would cost about \$350M, which is a reasonable figure. Both of these estimates are in the same range as the accounting estimates of \$200 per ton reported in Salvo (2010). This is fairly remarkable and a testament to the power of the MPNE framework given that these costs are inferred without any direct observation of investment expenditures in the data. I report the means and standard deviation of the differences in the last two columns. These costs are essentially unchanged across the two periods, which is in line with the a priori expectations that the Amendments only changed the sunk costs of entry during the period of time I observe.

The fixed and variable costs of divestment are estimated to be very large, in the

sense that firms will almost never find it reasonable to sell off significant amounts of their productive capacity. This reflects the paucity of downward substantial capacity adjustments observed in the data.³⁵ The numbers for before and after 1990 are very close, suggesting that the Amendments did not have a significant influence on divestment costs, as expected.

Finally, the distribution of exit costs is estimated to have a very low mean and fairly large standard deviation. This combination means that most firms will not find it worthwhile to exit unless they receive a very favorable draw from this distribution, as the estimated profits of remaining active firms are significantly positive even in the most contested markets. This is to be expected given the low exit rate of firms in this industry, particularly after 1990. Although the mean of the exit cost distribution shifts up after 1990, the standard deviation decreases. The combination of these two factors implies that exit continues to be a rare event after 1990. For example, the probability that a firm will receive a draw on the exit costs greater than \$75M is 3.4 percent before 1990 and 3.3 percent after 1990. The corresponding probabilities for draws greater than \$100M are 1.6 percent and 1.4 percent. These numbers are close to each other, which helps support the notion that the changes in the cost structure due to the Amendments primarily influenced entry costs. Further, the difference is not statistically significant.

The parameter on squared investment costs was set to zero after much experimentation. The reason is that the square tends to dominate the costs of investment for large changes; this implies that an entering firm building a reasonably large plant would face unreasonably large investment costs. This is an artifact of the fact that the quadratic adjustment costs are global, and most of the adjustments observed in the data do not span such large changes. I found the linear adjustment costs to give more reasonable results.

Distribution of Entry Costs I assume that the sunk costs of entry are independent draws from a normal distribution that is common across markets. I match the empirical probability of entry for a given state, given by the probit policy function, against the cumulative distribution function evaluated at the expected value of entering at that state. States were varied by the capacity of incumbent firms from 500,000

³⁵While there are a large number of reported divestments in the data, they typically are of small magnitude, and are often followed by positive investment of similar magnitude in the following period. I interpret these changes as most likely reflecting classical measurement error.

tons per year to 3M tons per year in 5,000 tons per year increments. The expected value of entry was computed using 250 replications at each state. The results of the estimation are presented in the bottom panel of Table 9.

One of the primary results of this paper is that I find the Amendments increased the sunk costs of entry. The mean of the entry cost distribution increased by 22 percent while the variance decreased by approximately 4 percent. The difference in means is statistically significant at standard levels while the difference in standard deviations is not. These two shifts work together to significantly decrease the chance of a firm receiving a small enough draw on the sunk cost of entry to warrant building a new facility. For example, the probability that a firm receives a draw on the entry costs below \$10M in the first period is about 4.5 percent; after the Amendments, the probability drops to 1.4 percent. If the threshold is raised to \$50M, the corresponding probabilities are 9.6 percent and 3.8 percent, respectively.³⁶ This is relevant because at the margin, the last entering firm can expect to make a relatively small amount of money in present value terms. Even when the firm is facing a large expected surplus conditional on entering, the shift in entry costs after 1990 greatly reduces the probability that a firm will find it optimal to do so. To emphasize, this change in the cost structure is the single most important determinant of the shift in market structure after 1990. As I show in the counterfactual simulations below, the increase in entry costs greatly reduces the chances that marginal firms enter a market, and this has significant effects on product market competition.

7 Policy Experiments

The benefit of estimating a structural model is the ability to simulate counterfactual policy experiments once a researcher knows the underlying primitives. My primary interest is to evaluate the welfare costs of the Amendments, so a natural investigation is to determine the differences across policy regimes for quantities of economic interest, including welfare measures for both producers and consumers. To achieve this, I compute the MPNE of the theoretical model with two sets of parameters: the observed post-Amendments cost structure, and the post-Amendments cost structure with the

³⁶If the variance of the post-1990 entry cost distribution is kept at the pre-1990 level while the mean shifts upward, the associated probabilities for draws less than \$10M and \$50M shift are 1.8 percent and 4.4 percent, respectively.

distribution of sunk entry costs taken from before the regulation.³⁷ It should be emphasized that these welfare calculations ignore the primary intended benefits of the Amendments: improved social welfare through cleaner air and its associated benefits. The results here should be interpreted in that light as a view on the changes in welfare due to changes in market structure as a result of the environmental regulation, and as such is only a part of the overall welfare changes of the 1990 Amendments to the Clean Air Act.³⁸

With policy functions from these equilibria it is possible to simulate hypothetical markets given some starting configuration. I compute the distribution of producer profits and consumer surplus under two different starting states: a new market with no incumbent firms and four potential entrants, and a market with two incumbents and two potential entrants. I take the parameters of demand for this market from Alabama, which is close to a representative market. Ideally, one would solve out for the equilibrium of every market in the US and simulate welfare changes for each one. Computational constraints, however, prevent this approach, and I have to restrict the number of active firms to be four, which is the median size of a cement market in the United States. While this is restrictive, the results with four firms indicate that the possibility of a fifth firm entering this market is very low. It is therefore reasonable to conclude that this restricted specification captures many of the essential dynamics of the median market.

Table 10 presents the results of the counterfactual simulations. The upper panel contains results for an initial state vector of four potential entrants and no incumbents. The lower panel reports results for an initial state vector of two incumbents, with joint capacities of 2.25M tons of cement per year, and two potential entrants. The first two columns report results using the distribution of entry costs from the pre-1990 period. The middle two columns report results using the distribution of entry costs from the post-1990 period. The last two columns report the difference between the

³⁷It is well known that these models potentially have many equilibria, some of which are not discoverable without sophisticated methods (see Doraszelski, et. al. (2008) for details). In the absence of any guarantees that the MPNE found here is unique, I report the solution method used to find the equilibrium so that my results will be reproducible. I used the b-spline interpolation methods described in Fowlie, Reguant, and Ryan (2010) to solve the resulting system. This process was stopped when the change in the norm of the value and policy functions was less than 1E-8—a sufficiently small level that changes from step to step only occurred in the fifth decimal place in the policy functions.

³⁸The integration of both views of the Amendments is an interesting research question that I am pursuing in Fowlie, Reguant, and Ryan (2010).

Table 10: Counterfactual Policy Experiments

	Low Entry Costs (Pre-1990)		High Entry Costs (Post-1990)		Difference	
	Mean	Standard Error	Mean	Standard Error	Mean	Standard Error
De Novo Market						
Total Producer Profit (\$ in NPV)	43,936.11	(7,796.98)	33,356.87	(7,767.22)	-11,182.04	(7,885.20)
Profit Firm 1 (\$ in NPV)	45,126.30	(10,304.87)	34,321.61	(9,520.93)	-11,965.22	(11,684.96)
Total Net Consumer Surplus (\$ in NPV)	1,928,985.09	(62,750.34)	1,848,872.52	(75,729.17)	-66,337.44	(58,404.32)
Total Welfare (\$ in NPV)	2,116,810.12	(74,265.74)	1,992,937.65	(96,634.83)	-119,771.39	(49,423.06)
Periods with no firms (periods)	1.29	(0.08)	1.32	(0.09)	0.04	(0.08)
Periods with one firm (periods)	1.51	(0.37)	2.60	(0.86)	1.05	(0.78)
Periods with two firms (periods)	8.17	(4.68)	21.43	(9.92)	12.26	(9.99)
Periods with three firms (periods)	54.71	(20.22)	91.35	(21.27)	33.38	(18.85)
Periods with four firms (periods)	135.91	(24.64)	84.03	(32.67)	-46.73	(25.04)
Average Size of Active Firm (tons)	980.71	(76.18)	1,054.65	(85.17)	73.42	(74.01)
Average Market Capacity (tons)	3,467.85	(188.21)	3,352.23	(208.94)	-112.75	(107.84)
Average Market Quantity (tons)	3,094.23	(161.57)	2,987.61	(177.58)	-105.69	(89.41)
Average Market Price	66.66	(1.90)	68.12	(2.11)	1.47	(1.14)
Mature Market						
Total Producer Profit (\$ in NPV)	223,292.75	(4,831.95)	231,568.23	(5,830.42)	9,551.01	(5,465.77)
Profit Firm 1 (\$ in NPV)	549,179.30	(14,138.37)	579,742.32	(20,446.75)	32,968.00	(19,161.33)
Total Net Consumer Surplus (\$ in NPV)	2,281,584.08	(52,663.88)	2,208,573.20	(62,906.14)	-62,974.37	(32,662.05)
Total Welfare (\$ in NPV)	3,178,504.60	(60,267.34)	3,141,916.43	(62,618.02)	-30,099.56	(18,078.21)
Periods with no firms (periods)	0.00	(0.00)	0.00	(0.00)	0.00	(0.00)
Periods with one firm (periods)	0.00	(0.00)	0.00	(0.00)	0.00	(0.00)
Periods with two firms (periods)	8.63	(3.57)	23.20	(10.05)	14.13	(10.00)
Periods with three firms (periods)	61.32	(16.83)	98.37	(21.49)	35.73	(20.16)
Periods with four firms (periods)	131.52	(20.10)	78.38	(31.99)	-50.00	(27.39)
Average Size of Active Firm (tons)	989.33	(44.45)	1,059.31	(63.41)	73.48	(54.64)
Average Market Capacity (tons)	3,502.49	(171.20)	3,371.03	(191.87)	-117.56	(73.19)
Average Market Quantity (tons)	3,123.42	(150.66)	3,001.98	(165.51)	-108.16	(69.48)
Average Market Price	66.82	(1.64)	68.36	(1.91)	1.44	(0.85)

Industry distributions were simulated along 25,000 paths of length 200 each. All values are present values denominated in thousands of dollars. The new market initially has no firms and four potential entrants. The incumbent market starts with one 750,000 TPY incumbent and one 1.5M TPY incumbent and two potential entrants. Counts of active firms may not sum to 200 due to rounding off. Means and standard deviations were calculated by subsampling.

two periods, with standard errors of the differences in parenthesis.

In the case of a new market, where the initial state vector is four empty slots waiting for entrants, overall welfare has decreased significantly due to the Amendments, declining by approximately \$120M (\$49M) in present value, a little over five percent. While artificial, the new market serves as a natural bound for the upper limit of welfare damages; as it is the market configuration that would be most affected by a change in sunk entry costs. Indeed, the driving factor for changes in welfare across both simulated markets is the change in entry rates. With the higher sunk costs of entry, the distribution of the number of active firms is shifted down. In equilibrium, the market is 68 percent (34 percent) more likely to have three active firms and 37 percent (18 percent) less likely to have four active firms under the higher entry costs. This compression of the firm distribution has significant effects on market outcomes. The average active firm is larger by about 7 percent (7 percent), reflecting the higher rate of return on investment given that firms expect softer competition in the future, which partially offsets the lower number of firms. Prices are 1.47 percent (1.14 percent) higher and quantities are 3.4 percent (2.8 percent) lower. However, the lower number of firms does not translate into better outcomes for producers, for whom profits decline from \$43.9M (\$7.7M) by \$11.1M (\$7.8M). This is an interesting result, as one may expect that higher costs would drive a wedge between the firms that actually enter and those that do not, increasing rents, but that is not the case in equilibrium. This is due to both the increased direct cost of entering the industry and the fact that firms are willing to enter at higher draws of entry costs knowing exactly that they will face fewer competitors in the future. In this sense, they compete away the potential projected oligopoly surplus induced by higher costs. On the other side of the market, consumer surplus decreases by 3.4 percent (3.0 percent) in this setting, a loss of \$66M (\$58M), which is 55 percent of the overall decline in total surplus.

The second market I consider has two incumbents with capacities of 750,000 TPY and 1.5M TPY. While the new entry market is an extreme case of what could have happened under the Amendments, a market with incumbents of over 2M TPY capacity is a close approximation of a mature, fully capitalized cement market of average size in the United States. As such, this should provide a lower bound to welfare penalties, as this market will be least affected by a change in entry rates. The primary effect, as in the new market, is the marked decrease in entry. The number of periods with four firms decreases by 40 percent (21 percent) under the post-1990 entry cost

distribution. The average size of an active firm is slightly larger in this case, again reflecting the fact that firms can recoup more of their investment costs with reduced product-market competition. However, this expansion in size is not sufficient to offset the reduced competitive effects of smaller numbers of firms in equilibrium: average prices increase by 1.44 percent (0.85 percent) and average quantities decrease by 3.3 percent (2.2 percent). As a result, consumer surplus is reduced by approximately \$62M (\$32M), a 2.8 percent decrease. On the other hand, producers enjoy a modest surplus increase of \$9.5M (\$5.5M), a 4.2 percent increase, under the higher entry costs. The profits of the largest firm increase even further, by \$32.9M (\$19.1M), or 6 percent, implying that incumbents experienced increases in profit at the expense of potential entrants, which highlights the within-industry distributional aspects of the legislation. As long as the costs of obtaining operating permits under Title V was lower than \$32.9M, the large incumbent is actually better off under the Amendments than before 1990. In this case, the static analysis of the engineering costs would not only ignore the dynamic costs to consumers, but also obtain welfare costs to suppliers of the wrong sign. The overall change in surplus in this market is a decrease of \$30M (\$18M), a decline of slightly less than 1.0 percent.

Extrapolating these costs to the entire US, under the assumption welfare losses can be summed equally across all 27 markets, leads to an estimate of over \$810M (\$486M) as a lower bound. The corresponding upper bound, \$3.2B (\$1.3B), clearly has little merit when extrapolated to the entire US, as it would be an estimate of the costs of starting the entire industry over from scratch under the two different sunk cost distributions. However, both numbers suggest that the welfare costs of the Amendments were significant, primarily through the reduction in product market competition. This result should be viewed carefully, however, as the reduction in output also reduces emissions in the short-run. In this sense, the negative consequences of environmental regulation through restricted competition in the product market are at least partially (and potentially more than) offset by reductions in emissions and their resulting welfare improvements.³⁹

One very strong assumption made here is that demand is not growing over time. It is difficult to assess what the effects of demand growth would be in this dynamic

³⁹Demonstrating the magnitude of these equilibrium effects is a complicated question beyond the scope of the present paper that I am pursuing in ongoing related research. See Fowlie, Reguant, and Ryan (2010) for more details.

setting. As seen with the new entrant market, strategic competition in entry and investment may undo some of the intuition regarding the effects of changing part of the cost structure. Growing demand makes the future more valuable relative to the world where it is not growing. One would expect this to increase the intensity of competition in entry, all else equal. On the other hand, it may be that firms that actually enter the market will do so at such a large size that this more than offsets the increased incentives for entry. This is an interesting case which I plan to explore in more detail in future research, as computational techniques and raw computing horsepower allow us to explore more complex state spaces than those considered here. In any case, the long-run effects may not be particularly pronounced in the United States, as many domestic cement markets appear to be relatively stable with respect to growth, as opposed areas of the world such as China where cement demand is booming.

8 Conclusion

In this paper, I have estimated the welfare costs of the 1990 Amendments to the Clean Air Act on the Portland cement industry. My principal finding is that a static analysis of the costs of the regulation will not only underestimate the costs to consumers, but will actually obtain estimates of the wrong sign for incumbent firms. Exploiting the timing structure of the implementation of the Amendments, I identify that the most significant economic change in the Portland cement industry was a large increase in the sunk costs of entry. As a result of lower entry rates, overall welfare decreased by at least \$810M. These results highlight the importance of estimating the welfare consequences of regulation using a dynamic model to account for all relevant changes to the determinants of market structure. A static model would also be incapable of calculating the counterfactual benefits to producers of paying higher entry costs but facing lower ex-post competition. The estimates that the certification process would at most cost \$5M per installation would underpredict the welfare costs by at least \$300M.

In Fowle, Reguant, and Ryan (2010), we extend of the analysis of the present paper is to examine the effects of a “cap-and-trade” market-based emissions control program, similar to the European Emissions Trading System for CO₂. In this environment the regulatory authority removes all specific point-source control requirements

and instead places an overall cap on the level of emissions. Firms are endowed with pollution rights which they are free to trade among each other. This type of policy has the benefit of achieving the most efficient configuration of production within the industry for a given level of pollution. However, it may have offsetting negative consequences by exacerbating the exercise of market power. There are a number of other interesting dynamic questions in this framework, from the nonlinear health effects of pollution concentration to the investment incentives of heterogeneous firms which we are pursuing in related work.

References

- Agency, U. S. E. P. (2001, June). NSR 90-day review background paper. Docket A-2001-19, Document II-A-01.
- Aguirregabiria, V. and P. Mira (2007). Sequential simulation-based estimation of dynamic discrete games. *Econometrica* 75(1), 1–53.
- Attanasio, O. (2000). Consumer durables and inertial behavior: Estimation and aggregation of (S, s) rules for automobile purchases. *Review of Economic Studies* 67, 667–96.
- Bajari, P., C. L. Benkard, and J. Levin (2007, September). Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–70.
- Bajari, P., H. Hong, and S. P. Ryan (2010). Identification and estimation of discrete games of complete information. Forthcoming in *Econometrica*.
- Benkard, C. L. (2004). A dynamic analysis of the market for widebodied commercial aircraft. *Review of Economic Studies* 71, 581–611.
- Bernheim, B. D. and M. D. Whinston (1990, Spring). Multimarket contact and collusive behavior. *RAND Journal of Economics* 21(1), 1–26.
- Besanko, D. and U. Doraszelski (2004, Spring). Capacity dynamics and endogenous asymmetries in firm size. *RAND Journal of Economics* 35(1), 23–49.
- Borkovsky, R., U. Doraszelski, and Y. Kryukov (2008). A user’s guide to solving dynamic stochastic games using the homotopy method. Forthcoming in *Operations Research*.
- Bureau, U. S. C. (1977–2002). Economic census.
- Chen, X. (2006). Large Sample Sieve Estimation of Semi-Nonparametric Models. *Handbook of Econometrics* 6.
- Chernozhukov, V. and H. Hong (2003). A MCMC approach to classical estimation. *Journal of Econometrics* 115(2), 293–346.

- Doraszelski, U. and M. Satterthwaite (2010, Summer). Computable markov-perfect industry dynamics. *RAND Journal of Economics* 41(2), 215–243.
- Ericson, R. and A. Pakes (1995). Markov perfect industry dynamics: A framework for empirical work. *Review of Economic Studies* 62(1), 53–82.
- Fershtman, C. and A. Pakes (2000). A dynamic game with collusion and price wars. *RAND Journal of Economics* 31(2), 207–36.
- Fowlie, M. (2009). Incomplete environmental regulation, imperfect competition, and emissions leakage. *American Economic Journal: Economic Policy* 1(2), 72–112.
- Fowlie, M., M. Reguant, and S. P. Ryan (2010). Pollution permits and the evolution of market structure. MIT Working Paper.
- Gowrisankaran, G. and R. Town (1997). Dynamic equilibrium in the hospital industry. *Journal of Economics and Management Strategy* 6(1), 45–74.
- Hall, G. and J. Rust (2000). The (S,s) rule is an optimal trading strategy in a class of commodity price speculation problems. Yale University Working Paper.
- Hotz, J., R. Miller, S. Sanders, and J. Smith (1994, April). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies* 61, 256–289.
- Jans, I. and D. I. Rosenbaum (1997). Multimarket contact and pricing: Evidence from the U.S. cement industry. *International Journal of Industrial Organization* 15, 391–412.
- Maskin, E. and J. Tirole (1988, May). A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica* 56, 549–69.
- Newey, W. K. and J. L. Powell (2003, September). Instrumental variable estimation of nonparametric models. *Econometrica* 71(5), 1565–1578.
- Pakes, A., M. Ostrovsky, and S. Berry (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics* 38(2), 373–399.
- Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* 75, 901–928.

- Politis, D. N. and J. P. Romano (1994). Large sample confidence regions based on subsamples under minimal assumptions. *The Annals of Statistics* 22(4), 2031–2050.
- Rothwell, G. and J. Rust (1995, December). Optimal response to a shift in regulatory regime: The case of the US nuclear power industry. *Journal of Applied Econometrics* 10, 79–122.
- Rust, J. (1987, September). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55(5), 999–1033.
- Ryan, S. P. and C. E. Tucker (2010). Diversification and the dynamics of technology adoption. Conditionally Accepted at *Quantitative Marketing and Economics*.
- Salvo, A. (2010). Inferring market power under the threat of entry: The case of the Brazilian cement industry. *RAND Journal of Economics* 41(2), 326–350.
- Scarf, H. E. (1959). The optimality of (s,S) policies in the dynamic inventory problem. In K. J. Arrow, S. Karlin, and P. Suppes (Eds.), *Chapter 13 in Mathematical Methods in the Social Sciences*. Stanford University Press.
- Tamer, E. (2002). Incomplete bivariate discrete response model with multiple equilibria. *Review of Economic Studies* 70, 147–167.
- van Oss, H. G. (2001, January). *Mineral Commodity Survey*. U.S. Geological Survey. Cement section.

A Identification

There are two sets of parameters in the model: those that are estimable without appeal to a dynamic model, and those that depend on the continuation value. The former category includes the demand curve and production costs, while the latter encompasses the costs of investment and divestment along with the distributions of fixed costs of investment, divestment, and exit.

The demand curve is nonparametrically identified under much weaker monotonicity and exclusion restrictions than imposed by the linear functional form in Equation

10.⁴⁰ The parameters of the production function are identified by functional form.⁴¹ The solution to the capacity-constrained Cournot game is unique, as the best-response curves are downward-sloping in their rivals' production. As the residual demand curve facing an individual firm moves in and out, it traces out the marginal cost of production. As mentioned earlier, the fixed cost of production would be identified when the firm chooses not to produce anything in a given period. However, in the present data sample all firms produce in all periods, so this parameter must be normalized to zero.

With regard to the dynamic parameters of the model, I provide a novel constructive approach to showing identification of two-step estimators, and demonstrate that the necessary and sufficient identification conditions are met in the present model. I also show how to estimate and identify parameters of unknown distributions in the underlying dynamic game, such as the distribution of fixed adjustment costs, which extends the class of models previously considered in the literature. The identification conditions are easy to verify, and apply to a wide class of dynamic games.

The approach to identifying the dynamic parameters is constructive. With policy functions for firms in hand, the econometrician can construct the ex-ante value function for firm i at state \bar{s} :

$$W_i(\bar{s}) = E_{s|\sigma, \bar{s}} \left[\sum_{t=0}^{\infty} \beta^t \pi_i(\sigma(s_t)) \right] = E_{s|\sigma, \bar{s}} \left[\sum_{t=0}^{\infty} \beta^t (1 - \theta) \cdot \begin{pmatrix} \bar{\pi}_i(s_t) \\ \zeta_i(\sigma(s_t)) \end{pmatrix} \right], \quad (31)$$

where the last equality follows by the linearity of the unknown parameters in the payoff function defined by Equation 5, and $\zeta_i(\sigma(s_t))$ is a vector of expected actions undertaken as state s_t , such as investment. The notation $E_{s|\sigma, \bar{s}}$ represents the integration over all possible paths of the state space in the future, conditional on the policy function, σ , which contains the probabilities of discrete choices and the levels of continuous choices. For discrete choices, these probabilities reflect optimal cutoff thresholds in the firm's private shock for undertaking a given discrete action. Optimality in equilibrium demands that no firm finds it payoff-increasing to make changes to these thresholds or levels.⁴² This implies that the derivative of the ex-ante value

⁴⁰See Newey and Powell (2003) and references therein for a general treatment of identification and estimation in nonparametric instrumental variables models.

⁴¹Functional form is a sufficient but not necessary condition for the identification of these parameters; given the availability of firm-specific cost shifters (capacity), these parameters are identified under more general conditions.

⁴²To an outside observer, deviations to the optimal threshold change the probability that a firm undertakes an action.

function with respect to the j -th aspect (either the level of an action or the probabilities of undertaking two or more actions) of $\sigma(s)$ at a single point in the state space, \hat{s} , is:

$$\frac{\partial W_i(\bar{s})}{\partial \sigma_{ij}(\hat{s})} = \theta \cdot E_{s|\sigma, \bar{s}} \sum_{t=0}^{\infty} \left[\beta^t \frac{\partial \zeta_i(\sigma(\hat{s}_t))}{\partial \sigma_{ij}(\hat{s}_t)} \right] + E_{s|\hat{\sigma}, \bar{s}} \sum_{t=0}^{\infty} \beta^t \left[(1 - \theta) \cdot \begin{pmatrix} \bar{\pi}_i(s_t) \\ \zeta_i(\sigma(s_t)) \end{pmatrix} \right] = 0. \quad (32)$$

There are two terms in the derivative: the change in the value function accruing to changes in the profit function, holding the distribution over states constant, and the change in the value function accruing from changes in the distribution over states expected to be visited in the future, holding per-period payoffs constant. Equation 32 neatly summarizes the opposing marginal costs and benefits that firms face when making optimal decisions. For example, firms weigh the marginal cost of investment against the marginal increase in product market profits when making optimal investment decisions.

Since the unknown parameters enter linearly, one can group terms such that:

$$-E_{s|\hat{\sigma}, \bar{s}} \sum_{t=0}^{\infty} \beta^t \bar{\pi}_i(s_t) = \theta \cdot \left[E_{s|\sigma, \bar{s}} \sum_{t=0}^{\infty} \left[\beta^t \frac{\partial \zeta_i(\sigma(\hat{s}_t))}{\partial \sigma_{ij}(\hat{s}_t)} \right] + E_{s|\hat{\sigma}, \bar{s}} \sum_{t=0}^{\infty} \beta^t \zeta_i(\sigma(s_t)) \right], \quad (33)$$

or equivalently,

$$y(\bar{s}, \sigma_{ij}(\hat{s})) = \theta \cdot x(\bar{s}, \sigma_{ij}(\hat{s})). \quad (34)$$

One can evaluate Equation 34 at $k = \dim(\theta)$ different states for the same perturbation, several different perturbations at the same state, or some mix of the two. In either case, one can then stack the resulting set of equations into a vector, Y , and the corresponding elements on the right-hand side into a matrix X , resulting in:

$$Y = \theta \cdot X. \quad (35)$$

The identification of θ then follows from the standard uniqueness conditions for a solution to ordinary least squares: as long as X has full rank, then θ is identified.

It remains to show that the estimated truncated fixed cost functions, $\tilde{\gamma}_{1i}(p_i)$, $\tilde{\gamma}_{4i}(p_d)$, and $\tilde{\phi}_i(p_e)$, identify their associated fixed cost distributions. It is necessary and sufficient to establish that the distribution function is one-to-one with the trun-

cated fixed cost function. I illustrate the identification arguments with the case of fixed costs of investment. First, I make the following support assumption:

Assumption 4. *There exists a set of states s such that a.) $p_d(s) = 0$ for all $p_i(s) \in (0, 1)$ and b.) $p_i(s) = 0$ for all $p_d(s) \in (0, 1)$.*

Assumption 4 is a support assumption on the equilibrium probabilities. Analogous assumptions have been used in the games literature to simplify multiple-factor inference problems into a single-factor problem. For example, in Tamer (2002), similar support conditions allow simultaneous entry games to be simplified into single-agent decision problems.⁴³ Here, the assumption allows the econometrician to invert the probability of investment onto the distribution of fixed investment costs, without having to worry about the convolution of divestment costs.⁴⁴

For clarity of notation, denote the value of making an investment, divestment, and doing nothing as:

$$V_i^+(s; \gamma_{i1}) = \max_{x_i^* > 0} \left[-\gamma_{i1} - \gamma_2 x_i^* - \gamma_3 x_i^{*2} + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right],$$

$$V_i^-(s; \gamma_{i4}) = \max_{x_i^* < 0} \left[-\gamma_{i4} - \gamma_5 x_i^* - \gamma_6 x_i^{*2} + \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) dP(s_i + x^*, s'_{-i}; s, \sigma(s)) \right],$$

and

$$V_i^0(s) = \beta \int E_{\epsilon_i} V_i(s'; \sigma(s'), \theta, \epsilon_i) dP(s_i, s'_{-i}; s, \sigma(s)).$$

The probability that a firm invests is equal to the following joint probability:

$$p_i(s) = Pr(V_i^+(s; \gamma_{i1}) > V_i^0(s), V_i^+(s; \gamma_{i1}) > V_i^-(s; \gamma_{i4})). \quad (36)$$

⁴³Tamer requires payoff shifters to go to infinity to drive the equilibrium probability of one player to zero for an action; this allows the econometrician to look at relationship between covariates and outcomes for the other player in isolation. Assumption 4 has the same flavor: it assumes that there exist states of the world where the econometrician observes the probability of either investment or divestment as being equal to zero.

⁴⁴This assumption requires zero probabilities, which are technically violated in the present application due to unbounded support on the errors of the targets and bands; there is always an infinitesimally small probability of having a firm receive an arbitrarily large shock which would induce either investment or divestment. Practically speaking, however, this is not a concern since I have verified that the computer is incapable of resolving the infinitesimal positive probability of this occurrence from zero. From the perspective of the estimator, you would obtain exactly the same results using either true zeros or the arbitrarily tiny probabilities implied by the estimated investment policy function.

This probability depends on the continuation values for investment, divestment, and doing nothing; the draw of fixed costs of investment; and critically for identification, also the draws of fixed costs of divestment and scrap values. Assumption 4 simplifies this problem by ensuring that there exists a part of the state space where the probability of investment is positive while the probability of divestment is approximately zero, which implies $Pr(V_i^+(s; \gamma_{i1}) > V_i^-(s; \gamma_{i4})) = 1$. The probability of observing investment is simplified:

$$p_i(s) = Pr(V_i^+(s; \gamma_{i1}) > V_i^0(s), V_i^+(s; \gamma_{i1}) > V_i^-(s; \gamma_{i4})) \quad (37)$$

$$\approx Pr(V_i^+(s; \gamma_{i1}) > V_i^0(s)), \quad (38)$$

where the second line follows from the assumption that the distribution of fixed investment costs is independent of the distribution of fixed costs of divestment. Letting $d(s)$ represent the direct and opportunity costs of investment, we can relate this probability to the distribution of fixed investment costs:

$$p_i(s) = Pr(\gamma_1 \leq d(s)) = F_\gamma(d(s)). \quad (39)$$

Define the inverse of the distribution function as follows:

$$F_\gamma^{-1}(p_i(s)) = \inf\{x \in R : p_i(s) \leq F_\gamma(x)\}. \quad (40)$$

If F is strictly increasing, F^{-1} is unique; otherwise it is the smallest value x such that the inequality is satisfied. In either case, knowledge of the inverse function fully characterizes the distribution function. By the definition of conditional expected value:

$$\tilde{\gamma}_{1i}(p_i(s)) = E(\gamma_1 | \gamma_1 \leq F_\gamma^{-1}(p_i(s))) = \frac{1}{p_i(s)} \int_{-\infty}^{F_\gamma^{-1}(p_i(s))} x f_\gamma(x) dx. \quad (41)$$

Multiplying both sides by $p_i(s)$ and differentiating with respect to $p_i(s)$ results in:

$$\frac{d}{dp_i(s)} \tilde{\gamma}_{1i}(p_i(s)) p_i(s) = F_\gamma^{-1}(p_i(s)) f_\gamma(F_\gamma^{-1}(p_i(s))) \frac{dF_\gamma^{-1}(p_i(s))}{dp_i(s)}. \quad (42)$$

Applying the definition of a derivative of an inverse function:

$$\frac{dF_\gamma^{-1}(p_i(s))}{dp_i(s)} = \frac{1}{f_\gamma(F_\gamma^{-1}(p_i(s)))}, \quad (43)$$

and substituting into Equation 42 obtains:

$$\frac{d\tilde{\gamma}_{1i}(p_i(s))p_i(s)}{dp_i(s)} = F_\gamma^{-1}(p_i(s)), \quad (44)$$

which establishes that the inverse distribution function is one-to-one in $\tilde{\gamma}_{1i}(p_i(s))p_i(s)$. The desired identification result follows from the fact that the inverse distribution function completely characterizes the distribution function. The distribution function can be completely nonparametrically recovered by allowing the degree of the sieve estimator to grow as the sample size goes to infinity.⁴⁵ It is possible to show the identification of the distributions of divestment and sunk exit costs in an analogous fashion.

The identification of the distribution of sunk entry costs is analogous to identification of a single-agent probit. Restating Equation 29:

$$Pr(\text{Entry}; s) = Pr(\kappa_i + \gamma_{1i} \leq EV^e(s)) = \Phi(EV^e(s); \mu_\kappa + \mu_\gamma, \sigma_\kappa^2 + \sigma_\gamma^2), \quad (45)$$

where μ_κ and σ_κ^2 are the mean and variance of the distribution of entry costs, which is distributed normally with CDF Φ . The terms μ_γ and σ_γ^2 represent the random fixed costs of investment; they enter as indicated since the sum of two normally-distributed variables is also distributed normally with mean and variance equal to the sum of their respective components. The distribution of fixed costs of investment is known, as discussed above. The probability of entry is known perfectly, and is a continuous function of the state variables, while the expected value of entering the market, $EV^e(s)$, is fully known from the behavior of incumbent firms. Identification requires that there exist two states, s and s' , such that $EV^e(s) \neq EV^e(s')$, which would be satisfied, for example, by considering the entry of a monopolist into two markets with differing levels of demand.

The present paper meets the requirements for identification. It is straightforward to check the rank condition on X in Equation 35 for a set of deviations required to

⁴⁵See Chen (2006) for details.

identify the structural and reduced-form parameters in Equation 17; intuitively, non-linearity in the per-period payoff function traces out these parameters. The truncated expected values are also one-to-one in their underlying distributions, as there are several combinations of observed states where the probability of investment varies while the probability of divestment and exit asymptote to zero. As divestment is highly unlikely at all states, the distribution of exit costs can be recovered by then examining states where the probability of exit is positive.